Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

VOL-I

Table of Content:

Table of Content:

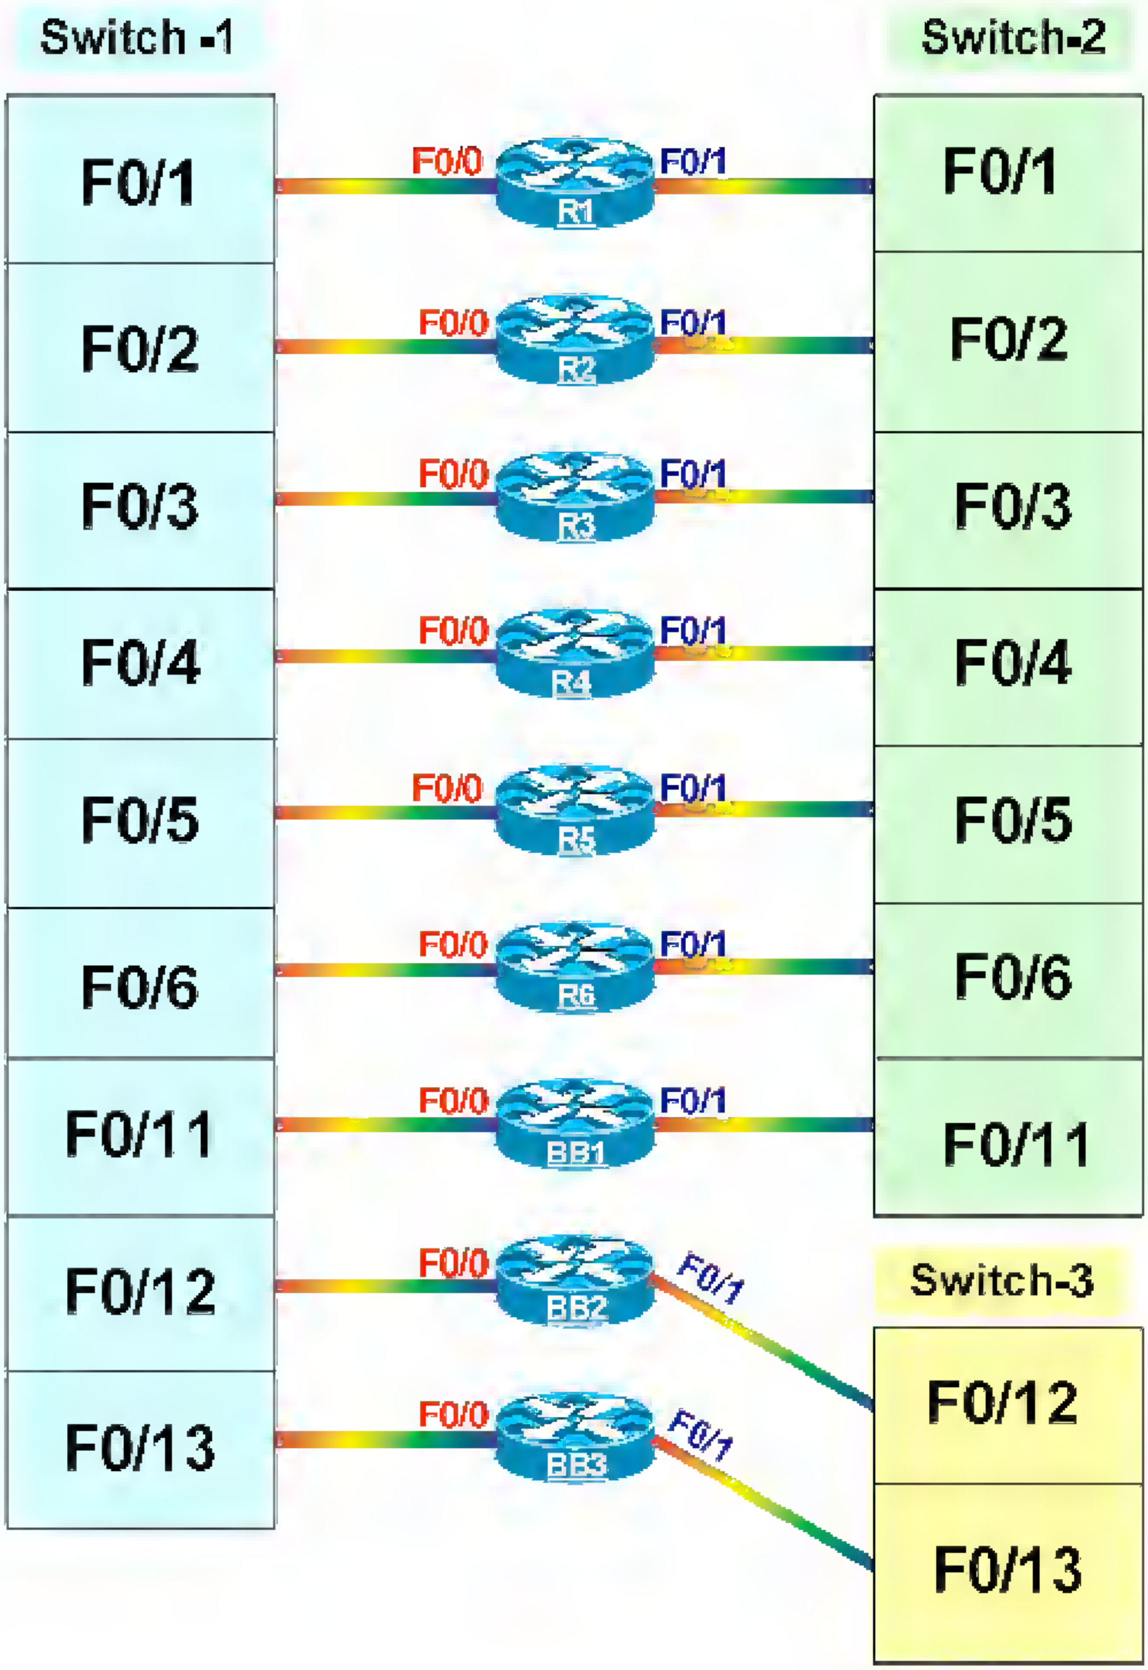
Switch -1
Switch-2
F0/1
F0/2
F0/3
F0/4
F0/5
F0/6
F0/11
F0/12
F0/13
F0/0
F0/1
R1
R2
R3
R4
R5
R6
BB1
BB2
BB3
Switch-3
F0/12
F0/13

The Serial connection between R1 and R3

The Serial connection between R4 and R5

Frame-relay Switch connections

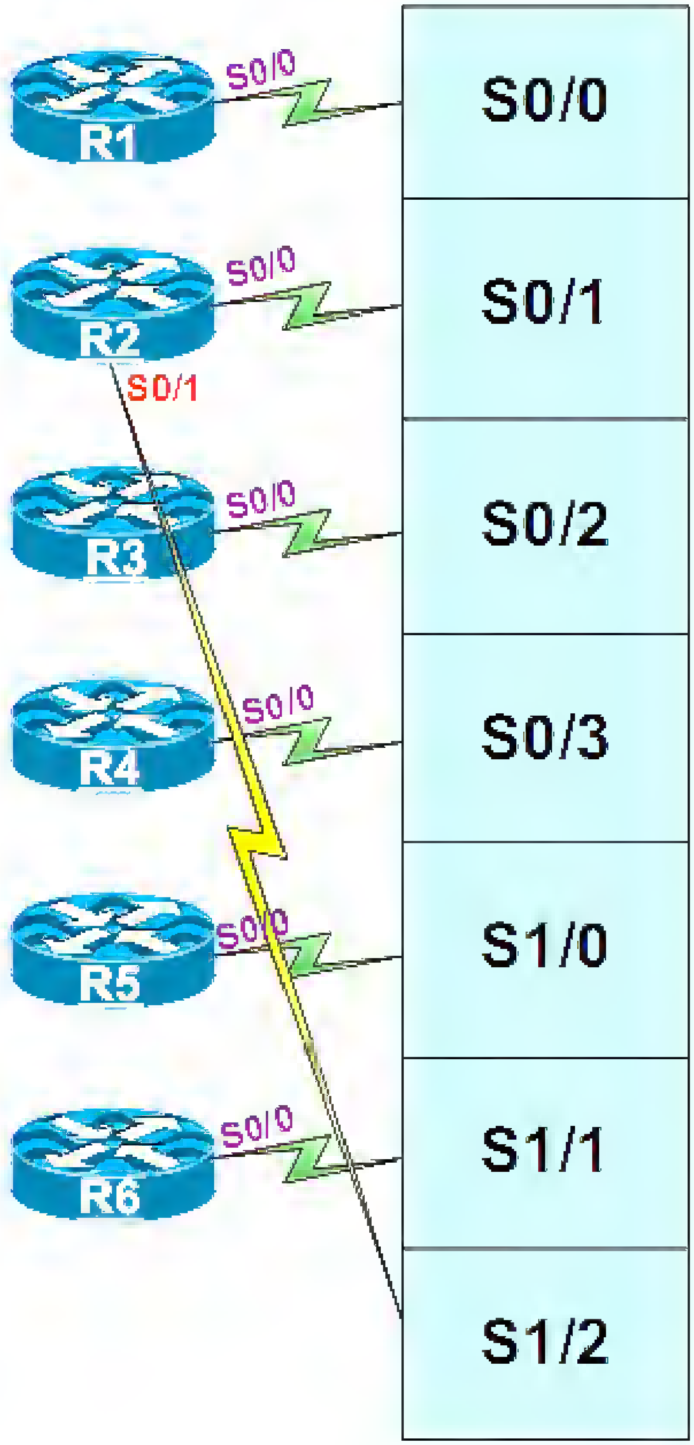

Frame-relay DLCI connections:

Router	Local DLCI	Connecting to:
R1	102	R2
	112	R2
	103	R3
	104	R4
	105	R5
	106	R6
R2	201	R1
	211	R1
	203	R3
	204	R4
	205	R5
	206	R6
R3	301	R1
	302	R2
	304	R4
	305	R5
	306	R6
R4	401	R1
	402	R2
	403	R3
	405	R5
	406	R6
R5	501	R1
	502	R2
	503	R3
	504	R4
	506	R6
R6	601	R1
	602	R2
	603	R3
	604	R4
	605	R5

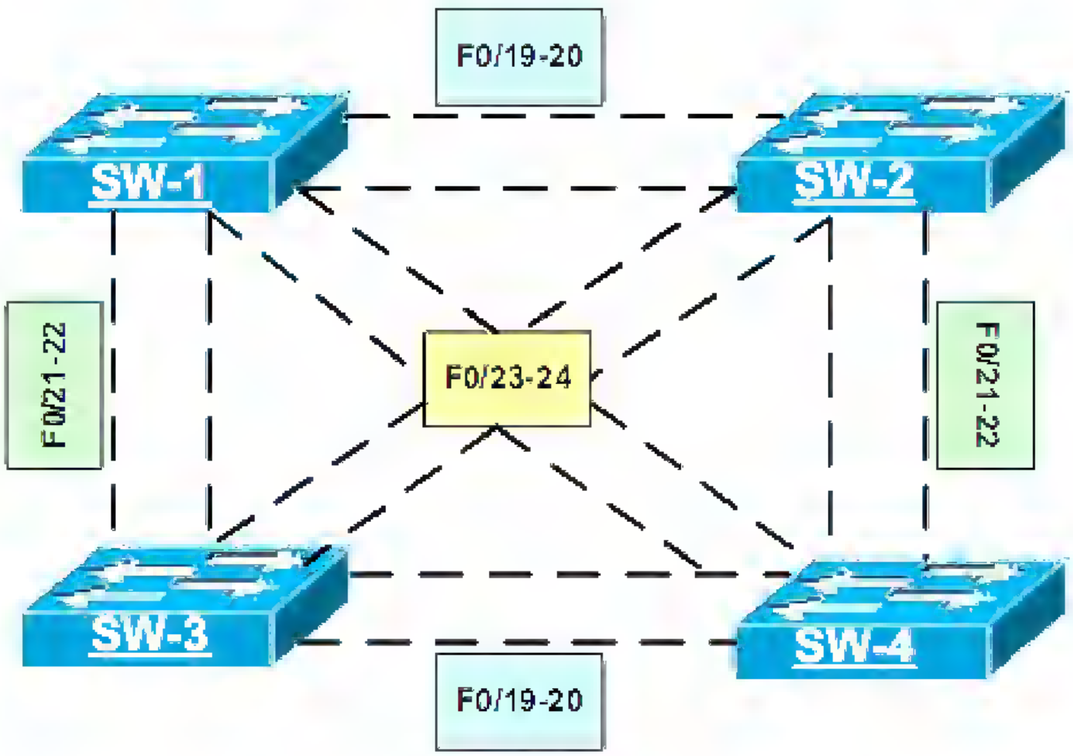
F0/19-20
SW-1
SW-2
F0/21-22
F0/23-24
F0/21-22
SW-3
SW-4
F0/19-20

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

3560
Switching

Lab 1
Trunks

Task 1

Shutdown all ports on all four switches and set the vtp domain name to TST.

On All Switches

(config)#int range f0/1 - 24
(config-if-range)#**Shut**

(config)#**vtp domain TST**

Task 2

Configure the following Hostnames:

The first Switch – Cat-1
The second Switch – Cat-2
The third Switch – Cat-3
The forth Switch – Cat-4

On the first Switch

Switch(config)#Hostname Cat-1

On the Second Switch

Switch(config)#Hostname Cat-2

On the Third Switch

Switch(config)#Hostname Cat-3

On the Forth Switch

Switch(config)#Hostname Cat-4

Task 3

Configure an ISL trunk between Cat-1 and Cat-2 using F0/19 interface based on the following policy:

Cat-1 - F0/19 → this port should be configured into permanent Trunking mode and it Should negotiate to convert the neighboring interface into a trunk

Cat-2 – F0/19 → this port should be configured to actively attempt to convert the link to A trunk

On SW1

Cat-1(config)#Int F0/19
Cat-1(config-if)#Switch mode Trunk

Note you get the following message:

Command rejected: An interface whose trunk encapsulation is "Auto" can not be configured to "trunk" mode.

The above message can be verified with the following show command:

Cat-1#Show interface F0/19 Switchport

Name: Fa0/19
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: On
(The rest of the output is omitted)

By default the ports on Catalyst 3560 are set to "Dynamic Auto" this is revealed by the "Administrative mode" and the Trunking encapsulation is set to "negotiate", revealed by "Administrative Trunking Encapsulation", when the "Administrative Trunking Encapsulation" is set to negotiate, the Trunking mode can NOT be set to ON.

To set the Trunking encapsulation to ISL:

On Cat-1

Cat-1(config)#Int F0/19
Cat-1(config-if)#**Switchport trunk encapsulation isl**

Cat-1(config-if)#No Shutdown

To verify the configuration:

On Cat-1

Cat-1#Show interface F0/19 Switchport

Name: Fa0/19

Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: isl
(The rest of the output is omitted)

To configure Cat-1

Cat-1(config)#int f0/19
Cat-1(config-if)#**Switchport mode trunk**

To verify the configuration:

Note the difference

On Cat-1

Cat-1#Show interface F0/19 Switchport

Name: Fa0/19
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: down
Administrative Trunking Encapsulation: isl
(The rest of the output is omitted)

Note the "Administrative Mode" is no longer "dynamic Auto" and the Trunking encapsulation is set to ISL.

On Cat-2

Cat-2(config)#int F0/19
Cat-2(config-if)#**Switchport mode dynamic desirable**
Cat-2(config-if)#No shut

To verify the configuration:

On Cat-2

Cat-2#Show interface F0/19 Switchport

Name: Fa0/19
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: Trunk
Administrative Trunking Encapsulation: **negotiate**
Operational Trunking Encapsulation: isl
(The rest of the output is omitted)

Note the operational mode changed from "Down" to "Trunk".

Note the mode is set to "Dynamic Desirable" and the "Administrative Trunking Encapsulation" is set to "negotiate" and the next line reveals the encapsulation mode that this port has negotiated, in this case ISL.

On Cat-1

Cat-1#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	isl	trunking	1

Port	Vlans allowed on trunk
Fa0/19	1-4094

Port	Vlans allowed and active in management domain
Fa0/19	1

Port	Vlans in spanning tree forwarding state and not pruned
Fa0/19	1

This column reveals the configured Trunking mode

Note Cat-2 negotiated an ISL Trunk, whereas, Cat-1 did not.

On Cat-2

Cat-2#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	desirable	n-isl	trunking	1

Port	Vlans allowed on trunk
Fa0/19	1-4094

```
Port        Vlans allowed and active in management domain
Fa0/19      1

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
```

Task 4

Configure an ISL trunk between Cat-1 and Cat-2 using F0/20 based on the following policy:

Cat-1 - F0/20 → this port should be configured into permanent Trunking mode and it Should negotiate to convert the neighboring interface into a trunk

Cat-2 – F0/20 → this port should be configured to negotiate a trunk ONLY if it receives Negotiate packets from a neighboring port; this port should never start The negotiation process

On Cat-1

```
Cat-1(config)#int f0/20
Cat-1(config-if)#Switchport trunk encap isl
Cat-1(config-if)#Switchport mode trunk
Cat-1(config-if)#NO shut
```

To verify the configuration:

On Cat-1

Cat-1#Show interface F0/20 Swi | inc Administrative Mode

Administrative Mode: **trunk**
(The rest of the output is omitted)

Cat-1#Sh inter status | inc Fa0/20

```
Fa0/20                notconnect    1          auto    auto    10/100BaseTX
```

Note just because the output states that this interface is in "notconnect" state, it does not mean that the interface is not connected to any device, it means that it has not detected any signaling from neighboring interface.

On Cat-2

Cat-2(config)#int f0/20
Cat-2(config-if)#**Switchport mode dynamic auto**
Cat-2(config-if)#**NO** shut

To verify the configuration:

On Cat-2

Cat-2#Show inter f0/20 Switchport | Inc **Administrative Mode**

Administrative Mode: **dynamic auto**
(The rest of the output is omitted)

Note the "Administrative Trunking Encapsulation" is set to "ISL" on Cat-1, whereas, on Cat-2 its set to "negotiate".
If this task stated that F0/20 on Cat-2 should negotiate ISL ONLY, then, configuring "switchport mode dynamic auto" will not suffice and the "Switchport trunk encapsulation isl" needs be added to the configuration of Cat-2's F0/20.

On Cat-1

Cat-1#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	isl	trunking	1
Fa0/20	on	**isl**	**trunking**	1

(The rest of the output is omitted)

On Cat-2

Cat-2#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	**desirable**	**n-isl**	**trunking**	1
Fa0/20	**auto**	**n-isl**	**trunking**	1

(The rest of the output is omitted)

Task 5

Configure an ISL Trunk between Cat-1 and Cat-3 using F0/21 interface. These ports should be configured to negotiate to convert the neighboring interface into an ISL trunk, but should NOT be in permanent trunking mode.

On Both Switches:

Cat-x(config)#int f0/21
Cat-x(config-if)#**switchport trunk encapsulation isl**

Cat-x(config-if)#**switchport mode dynamic desirable**
Cat-x(config-if)#**NO** shut

To verify the configuration:

On Cat-1

Cat-1#Show inter f0/21 switchport | **Inc Administrative Mode**

Administrative Mode: **dynamic desirable**

Cat-1#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	isl	trunking	1
Fa0/20	on	isl	trunking	1
Fa0/21	**desirable**	**isl**	**trunking**	**1**

(The rest of the output is omitted)

On Cat-3

Cat-3#Show inter f0/21 Switchport | **Inc Administrative Mode**

Administrative Mode: **dynamic desirable**

Cat-3#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/21	**desirable**	**isl**	**trunking**	**1**

(The rest of the output is omitted)

Task 6

Configure an ISL trunk between Cat-1 and Cat-3 using F0/22 interface based on the following policy:

Cat-1 – F0/22 → this port should be configured to actively attempt to convert the link to A trunk

Cat-3 – F0/22 → this port should be configured to negotiate a trunk ONLY if receives negotiation packets from a neighboring port; this port should never start the negotiation process

On Cat-1

Cat-1(config)#int f0/22
Cat-1(config-if)#**switchport trunk encapsulation isl**
Cat-1(config-if)#**swi mode dynamic desirable**
Cat-1(config-if)#NO shut

On Cat-3

Cat-3(config)#int f0/22
Cat-3(config-if)#**Switchport mode dynamic auto**
Cat-3(config-if)#NO shut

To verify the configuration:

On Cat-1

Cat-1#Show interface f0/22 Switchport | Inc Administrative Mode

Administrative Mode: **dynamic desirable**

Cat-1#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	isl	trunking	1
Fa0/20	on	isl	trunking	1
Fa0/21	desirable	isl	trunking	1
Fa0/22	desirable	isl	trunking	1

(The rest of the output is omitted)

On Cat-3

Cat-3#Show interface f0/22 switchport | Inc Administrative Mode

Administrative Mode: dynamic auto

Cat-3#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/21	desirable	isl	trunking	1
Fa0/22	auto	n-isl	trunking	1

(The rest of the output is omitted)

If the "Switchport trunk encapsulation ISL" was added to Cat-3's F0/22 interface, the "encapsulation" column in the output of the "Show interface trunk" command would have been "isl" and NOT "n-isl" which means negotiated ISL.

Task 7

Configure an ISL trunk between Cat-1 and Cat-4 using F0/23 interface; these switches should be configured into permanent trunking mode and negotiate to convert the neighboring interface into a trunk.

On Cat-1 & Cat-4

```
Cat-x(config)#int f0/23
Cat-x(config-if)#switchport trunk encapsulation isl
Cat-x(config-if)#Switchport mode trunk
Cat-x(config-if)#NO shut
```

To verify the configuration:

On Cat-1

Cat-1#Show inter F0/23 switchport | Inc Administrative Mode

Administrative Mode: trunk

Cat-1#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	isl	trunking	1
Fa0/20	on	isl	trunking	1

```
Fa0/21      desirable   isl          trunking     1
Fa0/22      desirable   isl          trunking     1
Fa0/23      on          isl          trunking     1
```
(The rest of the output is omitted)

On Cat-4

Cat-4#Sh int F0/23 swi | Inc Administrative Mode

Administrative Mode: **trunk**

Cat-4#Show inter trunk

```
Port       Mode      Encapsulation  Status       Native vlan
Fa0/23     on        isl            trunking     1
```
(The rest of the output is omitted)

Task 8

Configure an ISL trunk between Cat-1 and Cat-4 using interface F0/24; these ports should NOT use DTP to negotiate a Trunk.

On Cat-1

Cat-1(config)#int F0/24
Cat-1(config-if)#**switchport trunk encapsulation isl**
Cat-1(config-if)#**switchport mode trunk**
Cat-1(config-if)#**switchport nonegotiate**
Cat-1(config-if)#No shut

This command disabled DTP, but it MUST be configured after the "switchport mode trunk" command

On Cat-4

Cat-1(config)#int F0/24
Cat-1(config-if)#**switchport trunk encapsulation isl**
Cat-1(config-if)#**switchport mode trunk**
Cat-1(config-if)#**switchport nonegotiate**
Cat-1(config-if)#NO shut

To verify the configuration:

On Cat-1

Cat-4#Sh int F0/24 swi | Inc Administrative Mode|Negotiation

Administrative Mode: trunk
Negotiation of Trunking: Off

Cat-1#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	isl	trunking	1
Fa0/20	on	isl	trunking	1
Fa0/21	desirable	isl	trunking	1
Fa0/22	desirable	isl	trunking	1
Fa0/23	on	isl	trunking	1
Fa0/24	on	isl	trunking	1

(The rest of the output is omitted)

On Cat-4

Cat-4#Sh int F0/24 swi | Inc Administrative Mode|Negotiation

Administrative Mode: trunk
Negotiation of Trunking: Off

Cat-4#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/23	on	isl	trunking	1
Fa0/24	on	isl	trunking	1

(The rest of the output is omitted)

Task 9

Configure a Dot1q trunk between Cat-2 and Cat-4 using F0/21 interface based on the following policy:

Cat-2 - F0/21 → this port should be configured into a permanent Trunking mode and it Should negotiate to convert the neighboring interface into a trunk

Cat-4 – F0/21 → this port should be configured to actively attempt to convert the link to A trunk

On Cat-2

```
Cat-2(config)#Int F0/21
Cat-2(config-if)#Switchport trunk encapsulation dot1q
Cat-2(config-if)#Switch mode Trunk
Cat-2(config-if)#NO Shutdown
```

On Cat-4

```
Cat-4(config)#int f0/21
Cat-4(config-if)#switchport mode dynamic desirable
Cat-4(config-if)#NO shut
```

To verify the configuration:

On Cat-2

Cat-2#Sh int trunk | Exc isl

Port	Mode	Encapsulation	Status	Native vlan
Fa0/21	on	802.1q	trunking	1

(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | exc isl

Port	Mode	Encapsulation	Status	Native vlan
Fa0/21	desirable	n-802.1q	trunking	1

(The rest of the output is omitted)

Task 10

Configure a trunk between Cat-2 and Cat-4 using F0/22 interface; you should use an industry standard protocol for the trunking encapsulation based on the following policy:

Cat-2 - F0/22 → this port should be configured into permanent Trunking mode and it Should negotiate to convert the neighboring interface into a Trunk

Cat-4 – F0/22 → this port should be configured to negotiate a trunk ONLY if receives Negotiate packets from a neighboring port; this port should never start The negotiation process

On Cat-2

Cat-2(config)#int f0/22
Cat-2(config-if)#**Switchport trunk encap dot1q**
Cat-2(config-if)#**Switchport mode trunk**
Cat-2(config-if)#**NO** shut

On Cat-4

Cat-4(config)#int f0/22
Cat-4(config-if)#**swi mode dynamic auto**
Cat-4(config-if)#**NO** shut

To verify the configuration:

On Cat-2

Cat-2#Show int trunk | exc **isl**

Port	Mode	Encapsulation	Status	Native vlan
Fa0/21	on	802.1q	trunking	1
Fa0/22	**on**	**802.1q**	**trunking**	**1**

(The rest of the output is omitted)

On Cat-4

Cat-4#Sh int trunk | exc **isl**

Port	Mode	Encapsulation	Status	Native vlan
Fa0/21	desirable	n-802.1q	trunking	1
Fa0/22	**auto**	**n-802.1q**	**trunking**	**1**

(The rest of the output is omitted)

Task 11

Configure a Trunk link between Cat-3 and Cat-4 using F0/19 interface. These ports should be configured to negotiate to convert the neighboring interface into a dot1q trunk, but they should NOT be in permanent trunking mode.

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk encapsulation dot1q
Cat-x(config-if)#switchport mode dynamic desirable
Cat-x(config-if)#NO shut
```

To verify the configuration:

On Cat-3

Cat-3#sh int trunk | exc isl

```
Port      Mode        Encapsulation  Status      Native vlan
Fa0/19    desirable   802.1q         trunking    1
```

(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | Exc isl

```
Port      Mode        Encapsulation  Status      Native vlan
Fa0/19    desirable   802.1q         trunking    1
Fa0/21    desirable   n-802.1q       trunking    1
Fa0/22    auto        n-802.1q       trunking    1
```

(The rest of the output is omitted)

Task 12

Configure a Dot1q trunk between Cat-3 and Cat-4 using F0/20 interface based on the following policy:

Cat-3 – F0/20 → this port should be configured to actively attempt to convert the link to a Trunk. This port should NOT be in permanent trunking mode.

Cat-4 – F0/20 → this port should be configured to negotiate a trunk ONLY if receives Negotiation packets from a neighboring port; this port should never start the negotiation process.

On Cat-3

```
Cat-3(config)#int f0/20
Cat-3(config-if)#switchport trunk encapsulation dot1q
```

Cat-3(config-if)#swi mode dynamic desirable
Cat-3(config-if)#NO shut

On Cat-4

Cat-4(config)#int f0/20
Cat-4(config-if)#Switchport mode dynamic auto
Cat-4(config-if)#NO shut

To verify the configuration:

On Cat-3

Cat-3#Sh int trunk | Exc isl

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	desirable	802.1q	trunking	1
Fa0/20	desirable	802.1q	trunking	1

(The rest of the output is omitted)

On Cat-4

Cat-4#Sh int trunk | exc isl

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	desirable	802.1q	trunking	1
Fa0/20	auto	802.1q	trunking	1
Fa0/21	desirable	n-802.1q	trunking	1
Fa0/22	auto	n-802.1q	trunking	1

(The rest of the output is omitted)

Task 13

Configure a Dot1q trunk between Cat-2 and Cat-3 using F0/23 interface; these switches should be configured into permanent trunking mode and negotiate to convert the neighboring interface into a trunk.

On Both Switches:

Cat-x(config)#int F0/23

```
Cat-x(config-if)#switchport trunk encapsulation dot1q
Cat-x(config-if)#Switchport mode trunk
Cat-x(config-if)#NO shut
```

To verify the configuration:

On Cat-2

Cat-2#Sh int trunk | exc isl

```
Port      Mode       Encapsulation  Status      Native vlan
Fa0/21    on          802.1q        trunking     1
Fa0/22    on          802.1q        trunking     1
Fa0/23    on          802.1q        trunking     1
```

(The rest of the output is omitted)

On Cat-3

Cat-3#Sh int trunk | exc isl

```
Port      Mode       Encapsulation  Status      Native vlan
Fa0/19    desirable  802.1q         trunking     1
Fa0/20    desirable  802.1q         trunking     1
Fa0/23    on         802.1q         trunking     1
```

(The rest of the output is omitted)

Task 14

Configure a Dot1q trunk between Cat-2 and Cat-3 using interface F0/24; these ports should NOT use DTP to negotiate a Trunk.

On Both Switches:

```
Cat-x(config)#int F0/24
Cat-x(config-if)#Switchport trunk encapsulation dot1q
Cat-x(config-if)#Switchport mode trunk
Cat-x(config-if)#Switchport nonegotiate
Cat-x(config-if)#NO shut
```

To verify the configuration:

On Cat-2

Cat-2#Sh int trunk | exc isl

```
Port      Mode      Encapsulation  Status      Native vlan
Fa0/21    on        802.1q         trunking    1
Fa0/22    on        802.1q         trunking    1
Fa0/23    on        802.1q         trunking    1
Fa0/24    on        802.1q         trunking    1
```

(The rest of the output is omitted)

On Cat-3

Cat-3#Show int trunk | exc isl

```
Port      Mode         Encapsulation  Status      Native vlan
Fa0/19    desirable    802.1q         trunking    1
Fa0/20    desirable    802.1q         trunking    1
Fa0/23    on           802.1q         trunking    1
Fa0/24    on           802.1q         trunking    1
```

(The rest of the output is omitted)

Task 15

Configure the following VLANs on Cat-1 and ensure that they are propagated to the other switches:

VLANs 2 – 10, 100, 200, 300, 400, 120, 130, 140, 230, 240, and 340

On Cat-1

Cat-1(config)#vlan 2-10,100,200,300,400,120,130,140,230,240,340
Cat-1(config-vlan)#**exit**

To verify the configuration:

On All Switches:

Cat-x#Sh vlan brie | b VLAN0002

```
2    VLAN0002                 active
3    VLAN0003                 active
4    VLAN0004                 active
5    VLAN0005                 active
6    VLAN0006                 active
7    VLAN0007                 active
8    VLAN0008                 active
9    VLAN0009                 active
10   VLAN0010                 active
100  VLAN0100                 active
120  VLAN0120                 active
130  VLAN0130                 active
140  VLAN0140                 active
200  VLAN0200                 active
230  VLAN0230                 active
240  VLAN0240                 active
300  VLAN0300                 active
340  VLAN0340                 active
400  VLAN0400                 active
```
(The rest of the output is omitted)

Task 16

Configure the trunks based on the following policy:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/19	Cat-1 ←→ Cat-2	ONLY 120
2	F0/21	Cat-2 ←→ Cat-4	ONLY 240
3	F0/19	Cat-3 ←→ Cat-4	ONLY 340
4	F0/21	Cat-1 ←→ Cat-3	ONLY 130
5	F0/23	Cat-1 ←→ Cat-4	ONLY 140
6	F0/23	Cat-2 ←→ Cat-3	ONLY 230

Policy item 1: ←

The output of the following Show command reveals the default status of the trunk:

Cat-1#Show inter trunk | B Vlans allowed on trunk

Port   Vlans allowed on trunk

```
Fa0/19      1-4094
Fa0/20      1-4094
Fa0/21      1-4094
Fa0/22      1-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

To configure the task:

On Both Switches:

Note the following command ONLY allows VLAN 120 on the trunk

Cat-x(config)#int f0/19
Cat-x(config-if)#**Switchport trunk allowed VLAN 120**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      120
Fa0/20      1-4094
Fa0/21      1-4094
Fa0/22      1-4094
Fa0/23      1-4094
Fa0/24      1-4094
Port        Vlans allowed and active in management domain
Fa0/19      120
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
```

Fa0/24   1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)

Policy item 2:

On Cat-2 and Cat-4:

Cat-x(config)#int f0/21
Cat-x(config-if)#**switchport trunk allowed vlan 240**

To verify the configuration:

On Cat-4

Cat-2#Show int trunk | B Vlans allowed on trunk

Port   Vlans allowed on trunk
Fa0/19   120
Fa0/20   1-4094
Fa0/21   240
Fa0/22   1-4094
Fa0/23   1-4094
Fa0/24   1-4094

Port   Vlans allowed and active in management domain
Fa0/19   120
Fa0/20   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21   240
Fa0/22   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24   1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)

On Cat-2

Cat-4#Show int trunk | B Vlans allowed on trunk

Port   Vlans allowed on trunk
Fa0/19   1-4094
Fa0/20   1-4094
Fa0/21   240
Fa0/22   1-4094
Fa0/23   1-4094

```
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

Policy Item 3:

On Cat-3 and Cat-4

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk allowed vlan 340
```

To verify the configuration:

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      340
Fa0/20      1-4094
Fa0/21      1-4094
Fa0/22      1-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      340
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      340
Fa0/20      1-4094
Fa0/21      240
Fa0/22      1-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      340
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Policy Item 4:

On Cat-1 and Cat-3

Cat-x(config)#int f0/21
Cat-x(config-if)#**switchport trunk allowed vlan 130**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      120
Fa0/20      1-4094
Fa0/21      130
Fa0/22      1-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      120
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      130
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
```

Fa0/24   1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    340
Fa0/20    1-4094
Fa0/21    130
Fa0/22    1-4094
Fa0/23    1-4094
Fa0/24    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    340
Fa0/20    1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21    130
Fa0/22    1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23    1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24    1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Policy Item 5:

On Cat-1 and Cat-4

Cat-x(config)#int F0/23
Cat-x(config-if)#**Switchport trunk allowed vlan 140**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    120
Fa0/20    1-4094
Fa0/21    130
Fa0/22    1-4094
Fa0/23    140
```

```
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      120
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      130
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      140
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      340
Fa0/20      1-4094
Fa0/21      240
Fa0/22      1-4094
Fa0/23      140
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      340
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      140
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Policy Item 6:

On Cat-2 and Cat-3

Cat-x(config)#int F0/23
Cat-x(config-if)#**Switchport trunk allowed vlan 230**

To verify the configuration:

On Cat-2

Cat-2#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19       120
Fa0/20       1-4094
Fa0/21       240
Fa0/22       1-4094
Fa0/23       230
Fa0/24       1-4094

Port        Vlans allowed and active in management domain
Fa0/19       120
Fa0/20       1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21       240
Fa0/22       1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23       230
Fa0/24       1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19       340
Fa0/20       1-4094
Fa0/21       130
Fa0/22       1-4094
Fa0/23       230
Fa0/24       1-4094

Port        Vlans allowed and active in management domain
Fa0/19       340
Fa0/20       1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21       130
Fa0/22       1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23       230
Fa0/24       1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Task 17

Add VLANs to the allowed list of the trunk based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/19	Cat-1 ←→ Cat-2	100
2	F0/21	Cat-2 ←→ Cat-4	200
3	F0/19	Cat-3 ←→ Cat-4	300
4	F0/23	Cat-1 ←→ Cat-4	400

Policy Item 1:

On Cat-1 and Cat-2

Cat-x(config)#int f0/19
Cat-x(config-if)#**Switchport trunk allowed vlan add 100**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    100,120
Fa0/20    1-4094
Fa0/21    130
Fa0/22    1-4094
Fa0/23    140
Fa0/24    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    100,120
Fa0/20    1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21    130
Fa0/22    1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23    140
Fa0/24    1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Cat-2#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      1-4094
Fa0/21      240
Fa0/22      1-4094
Fa0/23      230
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      230
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Policy item 2:

On Cat-2 and Cat-4

```
Cat-4(config)#int f0/21
Cat-4(config-if)#Switchport trunk allowed vlan 200
```

To verify the configuration:

On Cat-2

Cat-2#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      1-4094
Fa0/21      200,240
Fa0/22      1-4094
Fa0/23      230
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      200,240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
```

```
Fa0/23      230
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      340
Fa0/20      1-4094
Fa0/21      200,240
Fa0/22      1-4094
Fa0/23      140
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      340
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      200,240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      140
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

Policy item 3:

On Cat-3 and Cat-4

```
Cat-x(config)#int f0/19
Cat-x(config-if)#Switchport trunk allowed vlan add 300
```

To verify the configuration:

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      300,340
Fa0/20      1-4094
Fa0/21      130
Fa0/22      1-4094
```

Fa0/23   230
Fa0/24   1-4094

Port   Vlans allowed and active in management domain
Fa0/19   300,340
Fa0/20   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21   130
Fa0/22   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23   230
Fa0/24   1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

Port   Vlans allowed on trunk
Fa0/19   300,340
Fa0/20   1-4094
Fa0/21   200,240
Fa0/22   1-4094
Fa0/23   140
Fa0/24   1-4094

Port   Vlans allowed and active in management domain
Fa0/19   300,340
Fa0/20   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21   200,240
Fa0/22   1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23   140
Fa0/24   1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)

Policy item 4:

On Cat-1 and Cat-4

Cat-x(config)#int F0/23
Cat-x(config-if)#**Switchport trunk allowed vlan add 400**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      1-4094
Fa0/21      130
Fa0/22      1-4094
Fa0/23      140,400
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      130
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      140,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      300,340
Fa0/20      1-4094
Fa0/21      200,240
Fa0/22      1-4094
Fa0/23      140,400
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      300,340
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      200,240
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/23      140,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Task 18

Remove VLANs from the allowed list of the trunks based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/22	Cat-1 ←→ Cat-3	Remove 1, 4 – 10 ONLY
2	F0/22	Cat-2 ←→ Cat-4	Remove 2, 4 – 10 ONLY

Policy item 1:

On Cat-1 and Cat-3

Cat-x(config)#int f0/22
Cat-x(config-if)#**Switchport trunk allowed vlan remove 1,4-10**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    100,120
Fa0/20    1-4094
Fa0/21    130
Fa0/22    2-3,11-4094
Fa0/23    140,400
Fa0/24    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    100,120
Fa0/20    1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21    130
Fa0/22    2-3,100,120,130,140,200,230,240,300,340,400
Fa0/23    140,400
Fa0/24    1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
```

```
Fa0/19     300,340
Fa0/20     1-4094
Fa0/21     130
Fa0/22     2-3,11-4094
Fa0/23     230
Fa0/24     1-4094

Port       Vlans allowed and active in management domain
Fa0/19     300,340
Fa0/20     1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21     130
Fa0/22     2-3,100,120,130,140,200,230,240,300,340,400
Fa0/23     230
Fa0/24     1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

Policy item 2:

On Cat-2 and Cat-4

Cat-x(config)#int f0/22
Cat-x(config-if)#**Switchport trunk allowed vlan remove 2,4-10**

To verify the configuration:

On Cat-2

Cat-2#Show int trunk | B Vlans allowed on trunk

```
Port       Vlans allowed on trunk
Fa0/19     100,120
Fa0/20     1-4094
Fa0/21     200,240
Fa0/22     1,3,11-4094
Fa0/23     230
Fa0/24     1-4094

Port       Vlans allowed and active in management domain
Fa0/19     100,120
Fa0/20     1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21     200,240
Fa0/22     1,3,100,120,130,140,200,230,240,300,340,400
Fa0/23     230
```

```
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      300,340
Fa0/20      1-4094
Fa0/21      200,240
Fa0/22      1,3,11-4094
Fa0/23      140,400
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      300,340
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/21      200,240
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
Fa0/23      140,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

Task 19

Configure Cat-1, Cat-2 and Cat-4 based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/20	Cat-1 ←→ Cat-2	None
2	F0/24	Cat-1 ←→ Cat-4	None

Policy Item #1

On Cat-1 and Cat-2

Cat-x(config)#int f0/20
Cat-x(config-if)#**Swi trunk allow vlan none**

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      none
Fa0/21      130
Fa0/22      2-3,11-4094
Fa0/23      140,400
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      none
Fa0/21      130
Fa0/22      2-3,100,120,130,140,200,230,240,300,340,400
Fa0/23      140,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

On Cat-2

Cat-2#Show int trunk | B Vlans allowed on trunk

```
Port        Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      none
Fa0/21      200,240
Fa0/22      1,3,11-4094
Fa0/23      230
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      none
Fa0/21      200,240
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
Fa0/23      230
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Policy Item #2

```
Cat-x(config)#int f0/24
Cat-x(config-if)#Swi trunk allowed vlan none
```

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      none
Fa0/21      130
Fa0/22      2-3,11-4094
Fa0/23      140,400
Fa0/24      none

Port      Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      none
Fa0/21      130
Fa0/22      2-3,100,120,130,140,200,230,240,300,340,400
Fa0/23      140,400
Fa0/24      none
```

(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19      300,340
Fa0/20      1-339,341-4094
Fa0/21      200,240
Fa0/22      1,3,11-4094
Fa0/23      140,400
Fa0/24      none

Port      Vlans allowed and active in management domain
Fa0/19      300,340
Fa0/20      1-10,100,120,130,140,200,230,240,300,400
Fa0/21      200,240
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
Fa0/23      140,400
```

Fa0/24 none
(The rest of the output is omitted)

Task 20

Configure Cat-1, Cat-3 and Cat-4 based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/20	Cat-3 ←→ Cat-4	All but 340
2	F0/22	Cat-1 ←→ Cat-3	All but 130

On Cat-3 and 4

Cat-x(config)#int f0/20
Cat-x(config-if)#**Swi trunk allowed vlan except 340**

To verify the configuration:

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    300,340
Fa0/20    1-339,341-4094
Fa0/21    130
Fa0/22    2-3,11-4094
Fa0/23    230
Fa0/24    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    300,340
Fa0/20    1-10,100,120,130,140,200,230,240,300,400
Fa0/21    130
Fa0/22    2-3,100,120,130,140,200,230,240,300,340,400
Fa0/23    230
Fa0/24    1-10,100,120,130,140,200,230,240,300,340,400
```
(The rest of the output is omitted)

On Cat-4

Cat-4#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    300,340
Fa0/20    1-339,341-4094
Fa0/21    200,240
Fa0/22    1,3,11-4094
Fa0/23    140,400
Fa0/24    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    300,340
Fa0/20    1-10,100,120,130,140,200,230,240,300,400
Fa0/21    200,240
Fa0/22    1,3,100,120,130,140,200,230,240,300,340,400
Fa0/23    140,400
Fa0/24    1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

On Cat-1 and Cat-3

Cat-x(config)#int f0/22
Cat-x(config-if)#**Swi trunk allowed vlan except 130**

To verify the configuration:

On Cat-1

Cat-1#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    100,120
Fa0/20    none
Fa0/21    130
Fa0/22    1-129,131-4094
Fa0/23    140,400
Fa0/24    none

Port      Vlans allowed and active in management domain
Fa0/19    100,120
Fa0/20    none
Fa0/21    130
Fa0/22    1-10,100,120,140,200,230,240,300,340,400
Fa0/23    140,400
```

Fa0/24 none
(The rest of the output is omitted)

On Cat-3

Cat-3#Show int trunk | B Vlans allowed on trunk

```
Port      Vlans allowed on trunk
Fa0/19    300,340
Fa0/20    1-339,341-4094
Fa0/21    130
Fa0/22    1-129,131-4094
Fa0/23    230
Fa0/24    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    300,340
Fa0/20    1-10,100,120,130,140,200,230,240,300,400
Fa0/21    130
Fa0/22    1-10,100,120,140,200,230,240,300,340,400
Fa0/23    230
Fa0/24    1-10,100,120,130,140,200,230,240,300,340,400
```

(The rest of the output is omitted)

Task 21

Configure Cat-2 and Cat-3 based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/23	Cat-2 ←→ Cat-3	ALL
2	F0/24	Cat-2 ←→ Cat-3	ALL

On Cat-2 and Cat-3

```
Cat-x(config)#int range f0/23-4
Cat-x(config-if)#swi trunk allow vlan all
```

To verify the configuration:

On Cat-2

```
Cat-2#Show int trunk | B Vlans allowed on trunk

Port        Vlans allowed on trunk
Fa0/19      100,120
Fa0/20      none
Fa0/21      200,240
Fa0/22      1,3,11-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      100,120
Fa0/20      none
Fa0/21      200,240
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)
```

On Cat-3

```
Cat-3#Show int trunk | B Vlans allowed on trunk

Port        Vlans allowed on trunk
Fa0/19      300,340
Fa0/20      1-339,341-4094
Fa0/21      130
Fa0/22      1-129,131-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      300,340
Fa0/20      1-10,100,120,130,140,200,230,240,300,400
Fa0/21      130
Fa0/22      1-10,100,120,140,200,230,240,300,340,400
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
(The rest of the output is omitted)
```

Task 22

Erase the config.text and Vlan.dat on all four switches and reload them before proceeding to the next task.

On All Four Switches

Cat-x#**Delete vlan.dat**

Cat-x#**Delete config.text**

Cat-x#**reload**

Task 23

Configure all four switches based on following requirements:

- Shut down all ports on all four switches
- Configure a Dot1q trunk between Switch 1 and 2 using port F0/19
- Set the VTP domain on Switch 1 and 2 to TST
- Name the first Switch to Cat-1 and the second Switch to Cat-2.

On The First Switch:

Switch(config)#Host Cat-1

On The Second Switch:

Switch(config)#Host Cat-2

On All Four Switches:

Cat-x(config)#int range f0/1-24
Cat-x(config-if-range)#Shut

On Cat-1 and Cat-2

Cat-x(config)#int F0/19
Cat-x(config-if)#swi trunk encapsulation dot1q
Cat-x(config-if)#swi mode trunk

```
Cat-x(config-if)#NO shut

Cat-x(config)#Vtp domain TST
```

Task 24

Configure VLAN 100 on Cat-1 and assign its F0/1 interface to this VLAN.

On Cat-1

```
Cat-1(config)#int f0/1
Cat-1(config-if)#Swi mode acc
Cat-1(config-if)#Swi acc v 100
Cat-1(config-if)#NO shut
```

To verify the configuration:

On Cat-1

```
Cat-1#Show vlan brie | Inc VLAN0100

100  VLAN0100                 active
```

Task 25

Configure the switches such that they restrict flooded traffic to those trunk links that the traffic must use to access the appropriate network device/s

This task is asking for VTP Pruning to be enabled, to understand VTP pruning, its helpful to know the VTP message types.

There are four types of VTP advertisements that are exchanged between the switches, and they are:

1. **Summary advertisements: An update sent by VTP servers or a client every 300 seconds or when a VLAN database change occurs. This update includes: VTP version, domain name, configuration revision number, time stamp, and number of subset advertisements.**

If the advertisement results from a VLAN database change, one or more subset advertisements will follow.

2. **Subset advertisements:** An update that follows a summary advertisement resulting from a change in the VLAN database. A subset advertisement includes the specific change/s that was made to a given VLAN/s.

3. **Advertisement requests from clients:** These are updates sent by a switch requesting more information so it can update its database. If and when a switch receives a VTP summary advertisement with a configuration revision number higher than its own, the local switch will send an advertisement request, requesting information about changes so it can update it's VLAN database. A switch operating in VTP server mode then responds with one or more subset advertisements.

4. **VLAN membership announcement:** These messages are generated by the switches when VTP Pruning is enabled and a port is associated to a given VLAN; these messages tell the neighboring switch that the local switch is interested in receiving traffic for that given VLAN. If the local switch does NOT send this message for a given VLAN, the neighboring switch will NOT send the traffic for that VLAN, and therefore the traffic for that VLAN will be pruned.

On Cat-1

Cat-1#Show interface pruning

Pruning not currently enabled in this device's VTP administrative domain.

Note the above message states that the pruning feature is NOT enabled. The output of the following messages reveals the same fact:

Cat-1#Show vtp status | Inc VTP Pruning Mode

VTP Pruning Mode : Disabled

To enable VTP Pruning:

Cat-1#Vtp Pruning

Pruning switched on

To verify the configuration:

On Cat-1

Cat-1#Show vtp status | Inc VTP Pruning Mode

VTP Pruning Mode : Enabled

Note this configuration will be propagated to all switches that have a trunk establishes with the local switch that and are in the same VTP domain:

On Cat-2

Cat-2#Show vtp status | Inc VTP Pruning Mode

VTP Pruning Mode : Enabled

Cat-2#Sh interface F0/19 pruning

Note the following output has two sections, the first section lists VLANs that are pruned, because the local switch has not received a Vlan Membership Announcement message (VMA) from the neighboring switch:

Port Vlans pruned for lack of request by neighbor
Fa0/19 none

This section of the output identifies for what VLANs the local switch has sent VMAs, and therefore, not pruned:

Port Vlan traffic requested of neighbor
Fa0/19 1

On Cat-1

Cat-1#Show interface f0/19 pruning

Note the local switch will NOT send traffic for VLAN out of this trunk interface, because the local switch has NOT received VMAs for this VLAN.

Port Vlans pruned for lack of request by neighbor
Fa0/19 100

Note the local switch has sent VMAs for these two VLANs:

Port Vlan traffic requested of neighbor
Fa0/19 1,100

Task 26

Configure VLANs 200, 300, 400, 500 and 600 on Cat-1 and ensure that these VLANs are propagated to Cat-2.

On Cat-1

```
Cat-1(config)#Vlan 200,300,400,500,600
Cat-1(config-vlan)#exit
```

On Cat-2

```
Cat-2#Show vlan br | exc unsup

<snip>
100  VLAN0100                    active
200  VLAN0200                    active
300  VLAN0300                    active
400  VLAN0400                    active
500  VLAN0500                    active
600  VLAN0600                    active
```

To verify the configuration:

On Cat-1

Note the output of the following show command displays that VLANs 100, 200, 300, 400, 500 and 600 are pruned:

```
Cat-1#Show interface F0/19 pruning

Port       Vlans pruned for lack of request by neighbor
Fa0/19     100,200,300,400,500,600

Port       Vlan traffic requested of neighbor
Fa0/19     1,100
```

On Cat-2

```
Cat-2#Show interface F0/19 pruning

Port       Vlans pruned for lack of request by neighbor
Fa0/19     200,300,400,500,600
```

Port Vlan traffic requested of neighbor
Fa0/19 1

Task 27

Configure F0/2 interface of Cat-2 in VLAN 100.

On Cat-2

```
Cat-2(config)#int f0/2
Cat-2(config-if)#swi mode acc
Cat-2(config-if)#swi acc v 100
Cat-2(config-if)#NO shut
```

Note you may have to wait for 30 seconds for convergence:

Cat-2#Show interface F0/19 pruning

```
Port      Vlans pruned for lack of request by neighbor
Fa0/19    200,300,400,500,600

Port      Vlan traffic requested of neighbor
Fa0/19    1,100
```

Note the output of the above show command reveals that the local switch has sent VMA message for VLAN 100.

Task 28

Configure the switches such that ONLY VLAN 300 is pruned.

On Cat-1

Cat-1#Show interface F0/19 pruning

```
Port      Vlans pruned for lack of request by neighbor
Fa0/19    200,300,400,500,600
```

Port   Vlan traffic requested of neighbor
Fa0/19   1,100

Note VLAN 300 is pruned. To configure the switches such that its no longer pruned:

On Both Switches:

Cat-x(config)#int f0/19
Cat-x(config-if)#**Switchport trunk pruning** vlan 300

Note the above command instructs the trunk to Prune VLAN 300 ONLY, therefore, the rest of the VLANs in the VLAN Database will NOT be pruned.

On Cat-1

Cat-1#Show interface F0/19 pruning

Port   Vlans pruned for lack of request by neighbor
Fa0/19   300

Port   Vlan traffic requested of neighbor
Fa0/19   1,100,200,400,500,600

Note VLAN 300 is the ONLY VLAN that is Pruned.

On Cat-2

Cat-2#Show interface F0/19 pruning

Port   Vlans pruned for lack of request by neighbor
Fa0/19   300

Port   Vlan traffic requested of neighbor
Fa0/19   1,100,200,400,500,600

Task 29

Configure the switches such that VLAN 200 is also pruned, you should NOT use the command from the previous task to accomplish this task.

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#Switchport trunk pruning vlan add 200
```

To verify the configuration:

On Cat-1

Cat-1#Sh inter f0/19 pruning

```
Port      Vlans pruned for lack of request by neighbor
Fa0/19    200,300

Port      Vlan traffic requested of neighbor
Fa0/19    1,100,400,500,600
```

On Cat-2

Note VLAN 200 is added to the list of Pruned VLANs

Cat-2#Show interface F0/19 pruning

```
Port      Vlans pruned for lack of request by neighbor
Fa0/19    200,300

Port      Vlan traffic requested of neighbor
Fa0/19    1,100,400,500,600
```

Task 30

Configure the switches such that NONE of the VLANs are pruned.

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#Switchport trunk pruning vlan NONE
```

To verify the configuration:

On Cat-1

Cat-1#Show interface f0/19 pruning

Port Vlans pruned for lack of request by neighbor
Fa0/19 none

Port Vlan traffic requested of neighbor
Fa0/19 1,100,200,300,400,500,600

On Cat-2

Note NONE of the VLANs are pruned

Cat-2#Show interface F0/19 pruning

Port Vlans pruned for lack of request by neighbor
Fa0/19 none

Port Vlan traffic requested of neighbor
Fa0/19 1,100,200,300,400,500,600

Task 31

Configure the switches such that all VLANs are pruned.

On Both Switches:

Cat-x(config)#Int F0/19
Cat-x(config-if)#**Switch trunk pruning vlan** 1,100,200,300,400,500,600

Note you should get the following errors:

Command rejected: Bad VLAN pruning list.

The reason the error message was generated was because **VLAN 1 CAN NOT BE PRUNED.**

Cat-x(config)#Int F0/19
Cat-x(config-if)#**Switch trunk pruning vlan** 100,200,300,400,500,600

To verify the configuration:

On Cat-1

Cat-1#Show interface F0/19 pruning

Port Vlans pruned for lack of request by neighbor
Fa0/19 200,300,400,500,600

Port Vlan traffic requested of neighbor
Fa0/19 1,100

Note VLAN 100 can NOT be pruned because the local switch has port membership in this VLAN.

On Cat-2

Cat-2#Show interface F0/19 pruning

Port Vlans pruned for lack of request by neighbor
Fa0/19 200,300,400,500,600

Port Vlan traffic requested of neighbor
Fa0/19 1,100

Task 32

Configure the switches such that VLAN 200 is no longer pruned; do not use a command that was used before to accomplish this task.

On Both Switches:

Cat-x(config)#int F0/19
Cat-x(config-if)#**Switchport trunk pruning vlan remove 200**

To verify the configuration:

On Cat-1

Cat-1#Show interface F0/19 pruning

Port Vlans pruned for lack of request by neighbor
Fa0/19 300,400,500,600

```
Port      Vlan traffic requested of neighbor
Fa0/19      1,100,200
```

Note VLAN 200 was removed from the list of VLANs being pruned.

On Cat-2

Cat-2#Show interface F0/19 pruning

```
Port      Vlans pruned for lack of request by neighbor
Fa0/19      300,400,500,600

Port      Vlan traffic requested of neighbor
Fa0/19      1,100,200
```

Task 33

Erase the vlan.dat and config.text and reload the switches before proceeding to the next lab.

Lab 2
EtherChannels

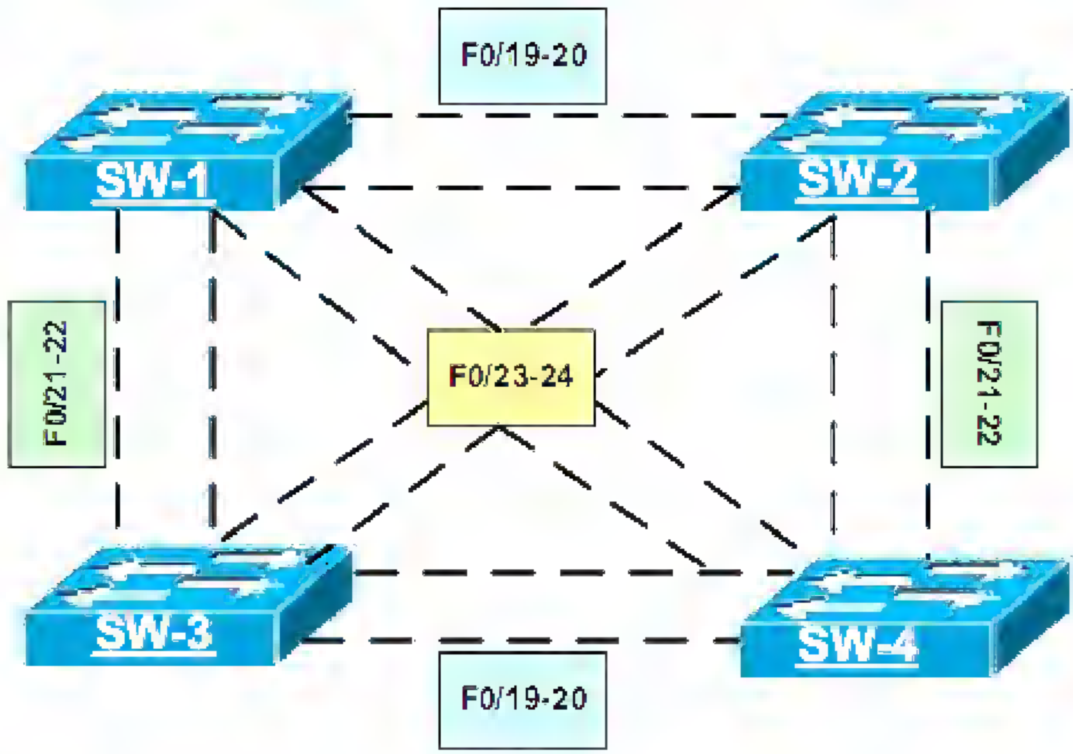

Task 1

Configure the hostname of the first switches as per diagram. Ensure that the ports of these four switches are in Shutdown mode. Configure VTP domain name to TST on all four switches.

On SW-1

```
Switch(config)#hostname SW-1

SW-1(config)#int range f0/1-24
SW-1(config-if-range)#Shut

SW-1(config)#VTP domain TST
```

On SW-2

```
Switch(config)#hostname SW-2

SW-2(config)#int range f0/1-24
SW-2(config-if-range)#Shut

SW-2(config)#VTP domain TST
```

On SW-3

```
Switch(config)#hostname SW-3

SW-3(config)#int range f0/1-24
SW-3(config-if-range)#Shut

SW-3(config)#VTP domain TST
```

On SW-4

```
Switch(config)#hostname SW-4

SW-4(config)#int range f0/1-24
SW-4(config-if-range)#Shut

SW-4(config)#VTP domain TST
```

Task 2

Configure ports F0/19 and F0/20 on SW-1 and SW-2 as trunk links using an industry standard protocol, these links should appear to Spanning-tree protocol as a single link. If one of the links fail, the traffic should use the other link without any interruption. The ports on SW-1 should be configured such that they ONLY respond to PAgP packets and never start the negotiation process.

EtherChannels provide the follows:

- Fault-tolerant, high speed links between switches and routers.
- EtherChannel provides an automatic recovery for the loss of a link by redistributing the traffic across the remaining link/s.
- STP will not block one of the links in the bundle because to STP, the bundle

looks like a single link.

- Up to 8 links can be combined to provide more bandwidth.
- The links within the bundle must have the same characteristics such as duplexing, speed and etc.
- EtherChannel can be configured as layer 2 or layer 3.
- With Layer 3, a logical interface (Port-Channel) is statically configured and all Layer 3 configurations are performed under that interface.
- With Layer 2, the logical interface is created automatically.
- With both Layer 2 and Layer 3, physical interfaces must be manually assigned to the logical interface using "channel-group" configuration command.
- EtherChannels can be configured automatically using Port aggregation protocol (PAgP) or Link Aggregation protocol (LACP).
- PAgP is a Cisco proprietary protocol, whereas LACP is an industry standard IEEE 802.3ad protocol.
- Switches can be configured to use PAgP by configuring them in AUTO or DESIRABLE mode.
- Switches can be configured to use LACP by configuring them in ACTIVE or PASSIVE mode.
- If the switches are configured in ON mode, they will not exchange LACP or PAgP packets.

There are 5 modes that the switches can be configured in:

- **ON** – Forces the interface into an EtherChannel without PAgP or LACP packets, both switches must be configured in ON mode for the EtherChannel to be established.

- **ACTIVE** – Used in LACP, the switches will actively negotiate an EtherChannel link.

- **PASSIVE** – Used in LACP, it places the interface in a passive negotiation mode where it only responds to LACP packets that it receives. In this mode the switch will not start the negotiation process; this setting minimizes the transmission of LACP packets.

- **AUTO** – Used in PAgP, it places the interface in a passive negotiation mode; It only responds to PAgP packets that it receives. In this mode the switch will not start the negotiation process; this setting minimizes the transmission of PAgP packets.

- **DESIRABLE** – Used in PAgP, the switches will actively negotiate an EtherChannel link.

The following table is very important to understand when configuring EtherChannels:

If SW-1 is configured in	If SW-2 is configured in	Will an EtherChannel be established?	The protocol used:
Desirable	Desirable	YES	PAgP
Desirable	Auto	YES	PAgP
Auto	Auto	NO	--
Active	Active	YES	LACP
Active	Passive	YES	LACP
Passive	Passive	NO	--
ON	ON	YES	NONE
ON	Auto	NO	--
ON	Desirable	NO	--
ON	Passive	NO	--
ON	Active	NO	--

When configuring EtherChannels, configuration of EtherChannels should be configured in certain order, the following is my recommendation for creating EtherChannels:

1. Configure "Default interface" for the interfaces involved.
2. Assign a channel-group and channel-group number to the physical interfaces, this step will create a port-channel interface automatically.
3. Configure the trunking encapsulation directly in port-channel interface configuration mode.
4. Reset the ports in the group by entering "Shut" and then, "No Shut".

Step One

On SW-1

SW-1(config)#**Default interface range F0/19-20**

SW-1(config)#Interface range f0/19-20
SW-1(config-if-range)#NO Shut

Step Two

SW-1(config)#int range f0/19-20
SW-1(config-if-range)#Channel-group 12 mode Auto

You should see the following messages:

Creating a port-channel interface Port-channel 12
Note the interface Port-channel 12 is created automatically:

SW-1#Show run | Inc interface Port-channel

interface Port-channel12

Step Three

SW-1(config)#**Int Port-channel 12**
SW-1(config-if)#**Switchport trunk encapsulation dot1q**
SW-1(config-if)#**Switchport mode trunk**

On SW-2

SW-2(config)#**Default interface range F0/19-20**

SW-2(config)#int range f0/19-20
SW-2(config-if-range)#**Channel-group 21 mode Desirable**

SW-2(config)#**Int Port-channel 21**
SW-2(config-if)#**Switchport trunk encapsulation dot1q**
SW-2(config-if)#**Switchport mode trunk**

Step Four

On SW-1 and SW-2

SW-x(config-if)#int range f0/19-20
SW-x(config-if-range)#**Shut**
SW-x(config-if-range)#**NO shut**

To verify the configuration:

On SW-1

SW-1#Sh interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Po12	on	802.1q	trunking	1

(The rest of the output is omitted)

On SW-2

SW-2#Sh interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Po21	on	802.1q	trunking	1

(The rest of the output is omitted)

On SW-1

SW-1#Show interface f0/19 switchport | Inc Operational Mode

Operational Mode: trunk (member of bundle Po12)

On SW-2

SW-2#Sh int f0/19 swi | Inc Operational Mode

Operational Mode: trunk (member of bundle Po21)

Task 3

Configure ports F0/21 and F0/22 on SW-3 and SW-1 as trunk links using an industry standard protocol, these links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should **NOT** negotiate by exchanging **LACP** or **PAgP** protocol to accomplish this task.

On SW-1

```
SW-1(config)#default interface range F0/21-22

SW-1(config)#Int range F0/21 - 22
SW-1(config-if-range)#Channel-group 13 mode on
SW-1(config-if-range)#NO shut

SW-1(config-if-range)#int port-channel 13
SW-1(config-if)#switchport trunk encapsulation dot1q
SW-1(config-if)#swi mode trunk
```

On SW-3

```
SW-3(config)#Default int range f0/21-22

SW-3(config)#Int range f0/21 – 22
```

```
SW-3(config-if-range)#Channel-group 31 mode on
SW-3(config-if-range)#NO shut

SW-3(config-if-range)#int port-channel 31
SW-3(config-if)#switchport trunk encapsulation dot1q
SW-3(config-if)#Swi mode trunk
```

On Both SW-1 and SW-3

```
SW-x(config)#int range f0/21-22
SW-x(config-if-range)#Shut
SW-x(config-if-range)#NO Shut
```

To verify the configuration:

On SW-1

SW-1#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Po12	on	802.1q	trunking	1
Po13	on	802.1q	trunking	1

(The rest of the output is omitted)

SW-1#Show etherchannel protocol

```
            Channel-group listing:
            ----------------------
Group: 12
---------
Protocol:  PAgP
```

Note PAgP is used for Etherchannel negotiation.

```
Group: 13
---------
Protocol:  -  (Mode ON)
```

Note PAgP or LACP is NOT in use

On SW-3

SW-3#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Po31	on	802.1q	trunking	1

(The rest of the output is omitted)

```
SW-3#Show etherchannel summary

Flags:  D - down          P - in port-channel
        I  - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3       S - Layer2
        U - in use        f - failed to allocate aggregator
         u - unsuitable for bundling
        w - waiting to be aggregated
         d - default port

Number of channel-groups in use: 1
Number of aggregators:         1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
31     Po31(SU)          -          Fa0/21(P)   Fa0/22(P)
```

Task 5

Ensure that all the EtherChannels created on SW-1 are load-balanced based on destination MAC address.

Etherchannel Load balancing can be done on 3550 or 3560 switches; 3560 switches are more flexible and provide more options, the following explains the load-balancing options available on these switches:

On 3550 Switches:

Source MAC Address – Packets forwarded to an EtherChannel are distributed across the ports in the channel based on the **Source MAC address** of the incoming packets. Therefore, different devices with different source MAC addresses use different interfaces in the bundle. When source MAC address load balancing is enabled, the load distribution based on the **Source and Destination IP address pair** is also enabled and this is for routed IP traffic.

Destination MAC Address – If the EtherChannel is between a router and a switch and since the router has a single MAC address, destination based load balancing is the best way. In this load balancing method, packets forwarded to an EtherChannel

are distributed across the ports in the channel based on the **Destination MAC address** of

the incoming packets.

Note there are only two choices on 3550 switches:

```
SW-3(config)#Port-channel load-balance ?
  dst-mac  Dst Mac Addr
  src-mac  Src Mac Addr
```

To verify the default setting:

On SW-3

SW-3#Show Etherchannel load-balance

EtherChannel Load-Balancing Operational State (src-mac):
Non-IP: Source MAC address
IPv4: Source MAC address
IPv6: Source IP address

SW-1 is a 3560:

The following are the options available on 3560 switches:

Source MAC Address – When packets are forwarded to an Etherchannel, they're distributed across the ports in the channel based on the **Source MAC address** of the incoming frame.

Destination MAC Address - When packets are forwarded to an Etherchannel, they're distributed across the ports in the channel based on the **Destination MAC address** of the incoming frame.

Source and Destination MAC Address - When packets are forwarded to an Etherchannel, they're distributed across the ports in the channel based on the **Source & Destination MAC address pair** of the incoming frame.

Source IP Address - When packets are forwarded to an Etherchannel, they're distributed across the ports in the channel based on the **Source IP address** of the incoming frame.

Destination IP Address - When packets are forwarded to an Etherchannel, they're distributed across the ports in the channel based on the **Destination IP address** of the incoming frame.

Source & Destination IP Address - When packets are forwarded to an Etherchannel,

they're distributed across the ports in the channel based on the **Source & Destination IP address pair** of the incoming frame.

To see the above options on 3560 switches:

```
SW-1(config)#Port-channel load-balance ?
  dst-ip        Dst IP Addr
  dst-mac       Dst Mac Addr
  src-dst-ip    Src XOR Dst IP Addr
  src-dst-mac   Src XOR Dst Mac Addr
  src-ip        Src IP Addr
  src-mac       Src Mac Addr
```

To verify the default setting:

```
SW-1#Show Etherchannel load-balance

EtherChannel Load-Balancing Operational State (src-mac):

Non-IP: Source MAC address
   IPv4: Source MAC address
   IPv6: Source IP address
```

To configure the load balancing based on the destination Mac addresses:

On SW-1:

```
SW-1(config)#port-channel load-balance dst-mac
```

To verify the configuration:

On SW-1

```
SW-1#Show etherchannel load

EtherChannel Load-Balancing Operational State (dst-mac):

Non-IP: Destination MAC address
    IPv4: Destination MAC address
    IPv6: Destination IP address
```

Note since the command is entered in the global configuration mode, it effects all EtherChannel ports created on the local switch.

Task 6

Ensure that all the EtherChannels created on SW-2 are load-balanced based on the following policy:

- For Non-IP, Source and Destination MAC address
- For IPv4, Source and Destination IP Address pair
- For IPv6, Source and Destination IP address pair

On SW-2

SW-2(config)#**port-channel load-balance src-dst-ip**

To verify the configuration:

On SW-2

SW-2#Show Etherchannel load-balance

EtherChannel Load-Balancing Operational State (src-dst-ip):
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination IP address
IPv6: Source XOR Destination IP address

The following reveals the behavior of a 3560 switch when the load balancing is changed:

If the load-balancing is changed to "src-mac":
Non-IP: Source MAC address
IPv4: Source MAC address
IPv6: Source IP address

If the load-balancing is changed to "dst-mac":
Non-IP: Destination MAC address
IPv4: Destination MAC address
IPv6: Destination IP address

If the load-balancing is changed to "src-ip":
Non-IP: Source MAC address

IPv4: Source IP address
IPv6: Source IP address

If the load-balancing is changed to "dst-ip":

Non-IP: Destination MAC address
IPv4: Destination IP address
IPv6: Destination IP address

If the load-balancing is changed to "src-dst-mac":
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination MAC address
IPv6: Source XOR Destination IP address

If the load-balancing is changed to "src-dst-ip":
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination IP address
IPv6: Source XOR Destination IP address

Task 7

Configure ports F0/21 and F0/22 on SW-2 and SW-4 as trunk links using Cisco proprietary trunking encapsulation, these links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should actively negotiate an etherchannel using PAgP.

On SW-2

SW-2(config)#**default interface range f0/21-22**

SW-2(config)#int range f0/21-22
SW-2(config-if-range)#**channel-group 24 mode desirable**
SW-2(config-if-range)#**NO** shut

SW-2(config)#Int port-channel 24
SW-2(config-if)#**switchport trunk encapsulation isl**
SW-2(config-if)#**switchport mode trunk**

On SW-4

SW-4(config)#**default interface range f0/21-22**

SW-4(config)#int range f0/21-22
SW-4(config-if-range)#**channel-group 42 mode desirable**
SW-4(config-if-range)#**NO** shut

SW-4(config)#int port-channel 42
SW-4(config-if)#switchport trunk encapsulation isl
SW-4(config-if)#switchport mode trunk

On SW-2 and SW-4

SW-4(config-if-range)#int range f0/21-22
SW-4(config-if-range)#shut
SW-4(config-if-range)#NO shut

To verify the configuration:

On SW-4

SW-4#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Po42	on	isl	trunking	1

(The rest of the output is omitted)

To verify the configuration:

On SW-4

SW-1#Show etherchannel protocol

Channel-group listing:

Group: 42

Protocol: PAgP

Note PAgP is used for Etherchannel negotiation.

On SW-2

SW-2#Show interface trunk

Port	Mode	Encapsulation	Status	Native vlan
Po21	on	802.1q	trunking	1
Po24	on	isl	trunking	1

(The rest of the output is omitted)

```
SW-2#Show etherchannel summary

Flags:  D - down         P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 2
Number of aggregators:           2

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
21     Po21(SU)        PAgP      Fa0/19(P)   Fa0/20(P)
24     Po24(SU)        PAgP      Fa0/21(P)   Fa0/22(P)
```

Task 8

Configure ports F0/19 and F0/20 on SW-3 and SW-4 as trunk links using Cisco proprietary trunking encapsulation, these links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports on SW-3 should be configured such that they ONLY respond to LACP packets that are received from the appropriate ports on SW-4.

On SW-3

```
SW-3(config)#default inter range f0/19-20

SW-3(config)#int range f0/19-20
SW-3(config-if-range)#channel-group 34 mode passive
SW-3(config-if-range)#NO shut

SW-3(config)#int port-channel 34
SW-3(config-if)#Switchport trunk encapsulation isl
SW-3(config-if)#Switchport mode trunk
```

On SW-4

```
SW-4(config)#default interface range f0/19-20

SW-4(config)#int range f0/19-20
SW-4(config-if-range)#channel-group 43 mode active
SW-4(config-if-range)#NO shut

SW-4(config)#int port-channel 43
SW-4(config-if)#switchport trunk encapsulation isl
SW-4(config-if)#switchport mode trunk
```

On SW-3 and SW-4

```
SW-4(config)#int range f0/19-20
SW-4(config-if-range)#Shut
SW-4(config-if-range)#NO shut
```

To verify the configuration:

On SW-3

```
SW-3#Show etherchannel protocol

                Channel-group listing:
                -----------------------
Group: 31
----------
Protocol:   -  (Mode ON)

Group: 34
----------
Protocol:  LACP
```

```
SW-3#Show interface trunk

Port      Mode         Encapsulation  Status        Native vlan
Po31      on           isl            trunking      1
Po34      on           isl            trunking      1
```

(The rest of the output is omitted)

On SW-4

```
SW-4#Show interface trunk
```

```
Port        Mode         Encapsulation  Status        Native vlan
Po42        desirable    n-isl          trunking      1
Po43        on           isl            trunking      1
```
(The rest of the output is omitted)

SW-4#Sh ether summ | B Number

Number of channel-groups in use: 2
Number of aggregators: 2

```
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
42      Po42(SU)      PAgP       Fa0/21(P)   Fa0/22(P)
43      Po43(SU)      LACP       Fa0/19(P)   Fa0/20(P)
```

Task 9

Configure ports F0/23 and F0/24 on SW-1 and SW-4 as trunk links using Cisco proprietary trunking encapsulation, these links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should be configured such that they actively negotiate a LACP Etherechannel.

On SW-1

```
SW-1(config)#default interface range f0/23-24

SW-1(config)#int range f0/23-24
SW-1(config-if-range)#channel-group 14 mode active
SW-1(config-if-range)#NO shut

SW-1(config)#int port-channel 14
SW-1(config-if)#switchport trunk encapsulation isl
SW-1(config-if)#switchport mode trunk
```

On SW-4

```
SW-4(config)#default interface range f0/23-24

SW-4(config)#int range f0/23-24
SW-4(config-if-range)#channel-group 41 mode active
SW-4(config-if-range)#NO shut
```

SW-4(config)#int port-channel 41
SW-4(config-if)#switchport trunk encapsulation isl
SW-4(config-if)#switchport mode trunk

On SW-1 and SW-4

SW-4(config-if)#int range f0/23-24
SW-4(config-if-range)#shut
SW-4(config-if-range)#NO shut

To verify the configuration:

On SW-1

SW-1#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Po12	on	802.1q	trunking	1
Po13	on	802.1q	trunking	1
Po14	on	isl	trunking	1

(The rest of the output is omitted)

On SW-4

SW-4#Show inter trunk

Port	Mode	Encapsulation	Status	Native vlan
Po41	on	isl	trunking	1
Po42	**desirable**	**n-isl**	**trunking**	1
Po43	on	isl	trunking	1

(The rest of the output is omitted)

SW-4#Show Etherchannel Pro

Channel-group listing:

Group: 41

Protocol: LACP

Group: 42

Protocol: PAgP

Group: 43

Protocol: LACP

Task 9

Configure ports F0/23 and F0/24 on SW-2 and SW-3 as a single layer three link; SW-2 should be configured with an IP address of 10.1.23.2 /24 and SW-3 should be configured with an IP address of 10.1.23.3 /24. These ports should NOT negotiate using LACP or PAgP.

Note when configuring layer 3 EtherChannels, I recommend the order of operation to be as follows:

1. **Default interface the physical interfaces**
2. **Configure the interface port-channel**
3. **Configure the port-channel interface with "NO Swi" and then configure the IP address**
4. **Configure the physical interfaces with "No Swi"**
5. **Assign the port-channel ID to the interfaces using the channel-group interface configuration command**
6. **Reset the physical interfaces by using "Shut" and "NO Shut"**

On SW-2

SW-2(config)#**default interface range f0/23-24**

SW-2(config)#int port-channel 23
SW-2(config-if)#**NO switchport**
SW-2(config-if)#**ip addr 10.1.23.2 255.255.255.0**

SW-2(config)#int range f0/23-24
SW-2(config-if-range)# **NO switchport**
SW-2(config-if-range)#**channel-group 23 mode on**
SW-2(config-if)#NO shut

On SW-3

SW-3(config)#**default interface range f0/23-24**

SW-3(config)#int port-channel 32

SW-3(config-if)#**NO switchport**
SW-3(config-if)#**ip addr 10.1.23.3 255.255.255.0**

SW-3(config)#int range f0/23-24
SW-3(config-if-range)#**Channel-group 32** mode on

Note if the "No Switchport" interface command is NOT configured, you should see the following error:

Command rejected (Port-channel32, Fa0/23): Either port is L2 and port-channel is L3, or vice-versa
% Range command terminated because it failed on FastEthernet0/23

SW-3(config-if-range)#NO swi
SW-3(config-if-range)#**Channel-group 32** mode on
SW-3(config-if-range)# NO shut

On SW-2 and SW-3

SW-3(config)#int range f0/23-24
SW-3(config-if-range)#Shut
SW-3(config-if-range)# NO shut

To verify and test the configuration:

On SW-2

SW-2#Show Etherchannel summary | B Number

Number of channel-groups in use: 3
Number of aggregators: 3

```
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
21     Po21(SU)       PAgP       Fa0/19(P)    Fa0/20(P)
23     Po23(RU)        -         Fa0/23(P)    Fa0/24(P)
24     Po24(SU)       PAgP       Fa0/21(P)    Fa0/22(P)
```

On SW-3

SW-3#Ping 10.1.23.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.23.2, timeout is 2 seconds:

```
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Task 10

Erase the startup configuration and vlan.dat before proceeding to the next lab

Lab 3
Basic 3560 configuration I

Task 1

Shutdown ports F0/21 – F0/24 on Switch 1 and 2.

On Both Switches

```
Switch(config)#int range F0/21-24
Switch(config-if-range)#Shut
```

To verify the configuration:

On Both Switches

Switch#Show int status

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/1		connected	1	auto	auto	10/100BaseTX
Fa0/2		connected	1	auto	auto	10/100BaseTX
Fa0/3		connected	1	a-full	a-100	10/100BaseTX
Fa0/4		connected	1	a-full	a-100	10/100BaseTX
Fa0/5		connected	1	a-full	a-100	10/100BaseTX
Fa0/6		connected	1	a-full	a-100	10/100BaseTX
Fa0/7		notconnect	1	auto	auto	10/100BaseTX
Fa0/8		notconnect	1	auto	auto	10/100BaseTX
Fa0/9		notconnect	1	auto	auto	10/100BaseTX
Fa0/10		connected	1	a-full	a-100	10/100BaseTX
Fa0/11		notconnect	1	auto	auto	10/100BaseTX
Fa0/12		notconnect	1	auto	auto	10/100BaseTX
Fa0/13		notconnect	1	auto	auto	10/100BaseTX
Fa0/14		notconnect	1	auto	auto	10/100BaseTX
Fa0/15		notconnect	1	auto	auto	10/100BaseTX
Fa0/16		notconnect	1	auto	auto	10/100BaseTX
Fa0/17		notconnect	1	auto	auto	10/100BaseTX
Fa0/18		notconnect	1	auto	auto	10/100BaseTX
Fa0/19		connected	1	a-full	a-100	10/100BaseTX
Fa0/20		connected	1	a-full	a-100	10/100BaseTX
Fa0/21		disabled	1	auto	auto	10/100BaseTX
Fa0/22		disabled	1	auto	auto	10/100BaseTX

```
Fa0/23                    disabled    1          auto   auto 10/100BaseTX
Fa0/24                    disabled    1          auto   auto 10/100BaseTX
```

(The rest of the output is omitted)

Task 2

Configure the first Switch to be in VTP domain called CCIE, this information should be propagated to Switch 2 via VTP messages. You can use any encapsulation or tagging to accomplish this task.

Before assigning a VTP domain name, there must be a trunk established between the two switches so the configurations will be propagated to the other switch.

On both switches

Switch#Show interface trunk

Switch#

Note the two 3560s switches are connected with 2 cross over ethernet cables, if these switches were 3550s, the two ports would have negotiated an ISL trunk, actually they would show up as "n-isl", this is because by default the ports are configured in desirable mode. With 3560 switches, the ports are not in desirable mode, a "show int f0/19 switchport" will reveal that by default the ports are configured in "Auto" mode (The Administrative Mode), and therefore, the port/s must be configured statically to trunk or negotiate a trunk.

On Both switches:

Switch#Show cdp neighbors

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme    Capability  Platform   Port ID
Switch           Fas 0/20          178         S I        WS-C3560-2Fas 0/20
Switch           Fas 0/19          177         S I        WS-C3560-2Fas 0/19
```

Note the "Show cdp neighbors" command reveals the ports connecting the two switches. The output may be different depending on the ports of the routers connecting to these switches; in this case the ports on the routers are in Shutdown mode.

On Both switches:

```
Switch(config)#int range f0/19-20
Switch(config-if-range)#switchport trunk encapsulation isl
Switch(config-if-range)#switchport mode trunk
```

To verify the configuration:

On the first switches:

```
Switch#Show int trunk

Port      Mode      Encapsulation  Status     Native vlan
Fa0/19    on        isl            trunking   1
Fa0/20    on        isl            trunking   1

Port      Vlans allowed on trunk
Fa0/19    1-4094
Fa0/20    1-4094

Port      Vlans allowed and active in management domain
Fa0/19    1
Fa0/20    1

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/19    1
Fa0/20    none
```

Now that the trunk is established between the two switches, therefore, the VTP configuration will be propagated via VTP messages:

On the first switch

```
Switch(config)#VTP domain CCIE
```

By default the 3560 switches are member of a domain called NULL, therefore, after entering the above command, you will get the following message unless the switch was member of another domain:

Changing VTP domain name from NULL to CCIE

This task could also be accomplished within the "VLAN database" as follows:

Switch#**Vlan database**
Switch(vlan)#**Vtp domain** CCIE
Switch(vlan)#Exit

When any configuration is performed in the Vlan database, you must configure the "exit" or the "apply" command for the changes to take effect.

Note the output of the following show command reveals that VTP propagated the VTP domain information to the second switch:

On the second switch:

Switch#Sh vtp status

```
VTP Version                         : 2
Configuration Revision              : 0
Maximum VLANs supported locally     : 1005
Number of existing VLANs            : 5
VTP Operating Mode                  : Server
VTP Domain Name                     : CCIE
VTP Pruning Mode                    : Disabled
VTP V2 Mode                         : Disabled
VTP Traps Generation                : Disabled
MD5 digest                          : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00
Local updater ID is 0.0.0.0 (no valid interface found)
```

Task 3

This VTP domain should be password protected using "**Cisco**" as the password.

On both switches

Switch(config)#**VTP password** Cisco

You should get the following message:

Setting device VLAN database password to Cisco

Note, if a domain name is not assigned to the switches and the default name of

"NULL" is used, a password can not be assigned.
The "VTP password" command can be entered in global configuration mode, privilege configuration mode or in the VLAN database mode.
The password command must be configured statically on both switches because this change will NOT get propagated via VTP messages.

To verify the configuration:

On the First switch

Switch#Show vtp status

```
VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 5
VTP Operating Mode              : Server     <------   The mode is server by default
VTP Domain Name                 : CCIE       <------   The domain name
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x14 0x7D 0x15 0x09 0xDC 0x39 0x65 0xC2
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00
Local updater ID is 0.0.0.0 (no valid interface found)
```

VTP password can be changed in three ways:

Privilege mode:

```
Switch#vtp password Cisco
```

Vlan Database:

```
Vlan database
  Vtp password Cisco
  Exit
```

Global config mode:

```
Switch(config)#vtp password Cisco
```

On the Second switch

Switch#Show vtp status

```
VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 5
VTP Operating Mode              : Server     <------   The mode is server by default
VTP Domain Name                 : CCIE       <------   The domain name
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Local updater ID is 0.0.0.0 (no valid interface found)
```

On any of the switches:

Switch#Show VTP password

This verifies the password, remember Spaces will not show

VTP Password: Cisco

Task 4

The first Catalyst switch should be configured with a hostname of Cat-1 and the second Catalyst should have a hostname of Cat-2.

On the first Switch

Switch(config)#Hostname Cat-1

On the Second Switch

Switch(config)#Hostname Cat-2

Task 5

Cat-2 should NOT have the ability to create, delete or rename VLAN or any VLAN information.

On Cat-2

Cat-2(config)#Vtp mode client

This configuration can be performed in the vlan database or global config mode. The above command displays the command as it was entered in the global config mode. If you are asked to enter the command in the vlan database, you must first enter the "vtp database" command in the privilege mode, then, enter "vtp client" and lastly the "exit" command is entered so the changes can take effect.

Once the command is entered you should get the following message:

Setting device to VTP CLIENT mode.

VTP Modes:

The switches can operate in three VTP modes and they are as follows:

- SERVER – The switch is able to delete, create, or rename VLAN information. Catalyst 3560 in server mode participates in the VTP domain and propagates the VLAN information.

- CLIENT – In this mode the switch is able to receive and process the VTP messages, but they are not able to create, delete, or rename VLAN information. They can assign a port to a given VLAN that already exists. Catalyst 3560 in client mode participates in the VTP domain and propagates the VTP messages.

- Transparent – In this mode the switch is able to create, delete and modify the VLAN information but it will not propagate its VLAN information to other switches. Catalyst 3560 switches in this mode do NOT participate in VTP domain. A Catalyst 3560 switch must be in this mode in order to create the extended-range VLANs (1006 – 4094), this configuration can only be performed in the global config mode and NOT in the Vlan database.

Task 6

Create and configure the following VLAN assignments on the switches:

Router Interface	VLAN number	CAT Switches Port
R1 – F0/0	12	Cat-1 / F0/1
R2 – F0/0	12	Cat-1 / F0/2
R3 – F0/0	34	Cat-1 / F0/3
R4 – F0/0	34	Cat-1 / F0/4
R5 – F0/0	56	Cat-1 / F0/5
R6 – F0/0	56	Cat-1 / F0/6

On Cat-1

```
Cat-1(config)#interface range f0/1 - 2
Cat-1(config-if)#switch mode access
Cat-1(config-if)#switch access vlan 12

Cat-1(config)#interface range f0/3 - 4
Cat-1(config-if)#switch mode access
Cat-1(config-if)#switch access vlan 34

Cat-1(config)#interface range F0/5 - 6
Cat-1(config-if)#switch mode access
Cat-1(config-if)#switch access vlan 56
```

Note the Vlan information will be propagated to the other switch (Cat-2), because both switches are in the same VTP domain and they are both configured with the same password.

On Cat-2

```
Cat-2#Show vlan brie | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2

12   VLAN0012                         active
34   VLAN0034                         active
56   VLAN0056                         active
```

Cat-2#Show VTP Status

```
VTP Version                          : 2
Configuration Revision               : 3
Maximum VLANs supported locally      : 1005
Number of existing VLANs             : 8
VTP Operating Mode                   : Client
VTP Domain Name                      : CCIE
VTP Pruning Mode                     : Disabled
VTP V2 Mode                          : Disabled
VTP Traps Generation                 : Disabled
MD5 digest                           : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC
0x1D
Configuration last modified by 0.0.0.0 at 3-1-93 00:06:11
Local updater ID is 0.0.0.0 (no valid interface found)
```

On Cat-1

Cat-1#Show VTP Status

```
VTP Version                          : 2
Configuration Revision               : 3
Maximum VLANs supported locally      : 1005
Number of existing VLANs             : 8
VTP Operating Mode                   : Server
VTP Domain Name                      : CCIE
VTP Pruning Mode                     : Disabled
VTP V2 Mode                          : Disabled
VTP Traps Generation                 : Disabled
MD5 digest                           : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC
0x1D
Configuration last modified by 0.0.0.0 at 3-1-93 00:06:11
Local updater ID is 0.0.0.0 (no valid interface found)
```

Note, the VTP version is 2, Configuration revision is 3, number of existing VLANs is 8 on both switches, (because they are synchronized), and the reason the VLAN information was propagated is because the VTP domain name and the password is identical on both switches and the switches are trunked.

Task 7

Configure Loopback 0 and Loopback 1 interfaces on Cat-1, use the IP address of 1.1.1.1 /8 and 11.1.1.1 /8 respectively and ensure that ONLY the IP address of Loopback 1 interface is used as the preferred source for the VTP IP updater address.

Note in the previous Task when the "show vtp status" command was entered on Cat-1, the last line of the output displayed "no valid interface found".
Catalyst switches will use the IP address of the lowest physical interface number, if one does not exist, then loopback 0 interface will be used as the source of all VTP messages, but this behavior can be change by using the "VTP interface Loopback 1" global config command.

On Cat-1

```
Cat-1(config)# Interface Loopback 0
Cat-1(config-if)# Ip address 1.1.1.1 255.0.0.0

Cat-1(config)# Interface Loopback 1
Cat-1(config-if)# Ip address 11.1.1.1 255.0.0.0
```

Cat-1#Show vtp status

```
VTP Version                      : 2
Configuration Revision           : 3
Maximum VLANs supported locally  : 1005
Number of existing VLANs         : 8
VTP Operating Mode               : Server
VTP Domain Name                  : CCIE
VTP Pruning Mode                 : Disabled
VTP V2 Mode                      : Disabled
VTP Traps Generation             : Disabled
MD5 digest                       : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC 0x1D
Configuration last modified by 0.0.0.0 at 3-1-93 00:06:11
Local updater ID is 1.1.1.1 on interface Lo0 (first layer3 interface found)
```

Note Loopback 0 is used as the source of all VTP messages. Enter the following command to change the source to Loopback 1 interface:

Cat-1(config)# **Vtp interface Loopback1 ONLY**

Note the "ONLY" argument makes this interface mandatory. YOU MUST TYPE LOOPBACK1 OR LO1, OR ELSE IT WILL NOT WORK, the IOS will take L1 but it WILL NOT WORK.

To verify the configuration:

On Cat-1

Cat-1#Show vtp status

```
VTP Version                      : 2
Configuration Revision           : 3
Maximum VLANs supported locally : 1005
Number of existing VLANs         : 8
VTP Operating Mode               : Server
VTP Domain Name                  : CCIE
VTP Pruning Mode                 : Disabled
VTP V2 Mode                      : Disabled
VTP Traps Generation             : Disabled
MD5 digest                       : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC 0x1D

Configuration last modified by 0.0.0.0 at 3-1-93 00:18:54
Local updater ID is 11.1.1.1 on interface Lo1 (preferred interface)
Preferred interface name is loopback1 (mandatory)
```

On Cat-2

Cat-2#Show vtp status

```
VTP Version                      : 2
Configuration Revision           : 3
Maximum VLANs supported locally : 1005
Number of existing VLANs         : 8
VTP Operating Mode               : Client
VTP Domain Name                  : CCIE
VTP Pruning Mode                 : Disabled
VTP V2 Mode                      : Disabled
VTP Traps Generation             : Disabled
MD5 digest                       : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC 0x1D
Configuration last modified by 0.0.0.0 at 3-1-93 00:22:29
```

Note this change has not been propagated, therefore, to force the propagation of this change, a VLAN is created, in this case VLAN 80, so you can see that the change was made by the Loopback 1 interface with an IP address of 11.1.1.1 on Cat-2. This VLAN should be deleted before proceeding to the next task.

On Cat-1

```
Cat-1(config)#Vlan 80
Cat-1(config-vlan)#Exit
```

To verify the configuration:

On Cat-2

Cat-2#Show vtp status

```
VTP Version                     : 2
Configuration Revision          : 4
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 9
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x02 0x05 0x92 0x34 0xF0 0xC0 0x35 0x9D
Configuration last modified by 11.1.1.1 at 3-1-93 00:34:33
```

On Cat-1

```
Cat-1(config)#No vlan 80
```

Task 8

Re-configure the trunk between the two switches such that none of these switches use DTP to negotiate this trunk.

On Both Switches

```
(config)#Interface range F0/19–20
(config-if-range)#Switchport nonegotiate
```

Note the ports must be in trunk mode before the "nonegotiate" command is entered, or else the following error message will be received:

Command rejected: Conflict between 'nonegotiate' and 'dynamic' status.

A port can be configured as follows:

Static Access – This port can belong to ONLY one VLAN, and it's manually assigned to a given VLAN.

Trunk – A trunk port by default is member of all normal range VLANs 1-1005 (but note that VLANs 1, 1002 – 1005 are automatically created and can not be removed, only 2 to 1001 can be manually created, these VLANs are kept in the VLAN.DAT).

This also includes the extended-range VLANs (1006 - 4094), and this membership can be limited by configuring the "allowed-vlan" command. This port can be encapsulated by ISL or tagged by 802.1q.

Dynamic Access – A dynamic access port can only be a member of one normal VLAN, and these ports are dynamically assigned to a given VLAN by a VMPS.

Voice VLAN – This is an access port connected to an IP phone such as Cisco's 7960, and this VLAN is used for Voice traffic.

Dot1q-Tunnel – These are tunnel ports and are used for 802.1q tunneling to

maintain customer VLAN integrity across a service provider's network. A tunnel port is configured on an edge switch in the service provider's network and it's connected to an 802.1q trunk port on a customer switch's interface, a tunnel port belongs to a single VLAN that is dedicated to tunneling.

To verify the configuration:

On Cat-1

Cat-1#Sh interfaces f0/19 switchport

Name: Fa0/19
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: isl
Operational Trunking Encapsulation: isl
Negotiation of Trunking: Off
(The rest of the output is omitted)

Cat-1#Sh interfaces f0/20 switchport

Name: Fa0/20

```
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: isl
Operational Trunking Encapsulation: isl
Negotiation of Trunking: Off
(The rest of the output is omitted)
```

Task 9

Configure the switches such that flooded traffic is restricted to the trunk links that the traffic must use to reach the destination device.

To see the default setting:

On Cat-2

Cat-2#Show vtp status

```
VTP Version                        : 2
Configuration Revision             : 5
Maximum VLANs supported locally    : 1005
Number of existing VLANs           : 8
VTP Operating Mode                 : Client
VTP Domain Name                    : CCIE
VTP Pruning Mode                   : Disabled  Pruning is disabled
VTP V2 Mode                        : Disabled
VTP Traps Generation               : Disabled
MD5 digest                         : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC 0x1D
Configuration last modified by 11.1.1.1 at 3-1-93 00:12:48
```

Note VTP Pruning is disabled by default, enter the following command to enable VTP pruning:

On Cat-1

Cat-1#**Vtp pruning**

This command can be configured in privilege mode, Global config mode, and/or in the Vlan database. Once this feature is enabled it will get propagated to the other

switches within the VTP domain.

To verify the configuration on both switches:

On Cat-2

Cat-2#Show vtp status

```
VTP Version                     : 2
Configuration Revision          : 5
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Enabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x97 0x9D 0xF1 0xF9 0xFE 0x21 0xCC 0x1D
Configuration last modified by 11.1.1.1 at 3-1-93 00:12:48
```

Note VTP messages propagate the change through the entire VTP domain.

Task 10

Configure Cat-1 and Cat-2 such that only the trunk ports (F0/19 and F0/20) and the ports that routers R1 to R6 are connected are in use, the rest of the ports should be configured in administratively down state.

On Both Switches:

```
(config)#Int range f0/7-18 , F0/23-24
(config-if-range)#Shut
```

To verify the configuration:

On Cat-1

Cat-1#Show inter status | Inc disable

```
Fa0/7                 disabled     1          auto    auto 10/100BaseTX
```

```
Fa0/8            disabled   1      auto   auto 10/100BaseTX
Fa0/9            disabled   1      auto   auto 10/100BaseTX
Fa0/10           disabled   1      auto   auto 10/100BaseTX
Fa0/11           disabled   1      auto   auto 10/100BaseTX
Fa0/12           disabled   1      auto   auto 10/100BaseTX
Fa0/13           disabled   1      auto   auto 10/100BaseTX
Fa0/14           disabled   1      auto   auto 10/100BaseTX
Fa0/15           disabled   1      auto   auto 10/100BaseTX
Fa0/16           disabled   1      auto   auto 10/100BaseTX
Fa0/17           disabled   1      auto   auto 10/100BaseTX
Fa0/18           disabled   1      auto   auto 10/100BaseTX
Fa0/21           disabled   1      auto   auto 10/100BaseTX
Fa0/22           disabled   1      auto   auto 10/100BaseTX
Fa0/23           disabled   1      auto   auto 10/100BaseTX
Fa0/24           disabled   1      auto   auto 10/100BaseTX
```

Task 11

Ensure that Cat-1 is the root bridge for the VLANs 12, 34 and Cat-2 is the root bridge for VLAN 56. Do NOT use the "priority" command to accomplish this task.

There are two commands that can be used to display the BID for a given switch:

- **Show version**
- **Show spanning-tree bridge**

On Cat-1

Cat-1#Show version | Inc Base

```
Base ethernet MAC Address       : 00:1B:D4:59:A6:00
```

The following command reveals the base MAC address of the switch; The BID is a combination of priority and the base MAC address.

Cat-1#Show spanning-tree bridge

```
                                                  Hello  Max  Fwd
Vlan                        Bridge ID             Time   Age  Dly  Protocol
---------------- --------------------------------- -----  ---- ---  --------
VLAN0001         32769 (32768,  1) 001b.d459.a600  2      20   15   ieee
VLAN0012         32780 (32768, 12) 001b.d459.a600  2      20   15   ieee
```

```
VLAN0034      32802 (32768, 34) 001b.d459.a600  2   20   15   ieee
VLAN0056      32824 (32768, 56) 001b.d459.a600  2   20   15   ieee
```

Note the priority starts with 32768, each VLAN that is created adds it's VLAN number to the default priority value (If the base priority and the VLAN number is added within the parenthesis, the sum will be the priority for that given VLAN), VLAN 12 adds 12 to the default priority value therefore the priority is 32780 and VLAN 34 adds 34 to the default priority value, therefore, the priority is 32802. Note that the MAC is the base MAC address and it remains the same, in this case (001b.d459.a600).
Note your MAC address maybe different.

Enter the following command to reveal the BID and the root bridge for a given VLAN:

On Cat-1

Cat-1#Show spanning-tree vlan 12

```
VLAN0012
  Spanning tree enabled protocol ieee          The MAC address of the root bridge
  Root ID    Priority    32780
             Address     0011.bbeb.8780
             Cost        19
             Port        21 (FastEthernet0/19)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32780 (priority 32768 sys-id-ext 12)
             Address     001b.d459.a600
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
                                               The Mac address of the local switch
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Root FWD 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
```

Enter the following commands to configure Cat-1 to be the root bridge for VLANs 12 and 34:

On Cat-1

Cat-1(config)#**Spanning-tree vlan 12,34 root primary**

The above command configures Cat-1 to be the root for VLANs 12 and 34; the "root" keyword is a macro that reduces the BID of the switch for a given VLAN by a value of 8192

(The lower value is the preferred value). There are no spaces between the 12 and the comma and the 34.

Cat-1#Show spanning-tree vlan 12

```
VLAN0012                              Note 32768+12-8192 = 24588
  Spanning tree enabled protocol ieee
  Root ID    Priority    24588
             Address     001b.d459.a600
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24588  (priority 24576 sys-id-ext 12)
             Address     001b.d459.a600
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
```

On Cat-2

Cat-2(config)##**Spanning-tree vlan 56 root primary**

To verify the configuration:

On Cat-2

Cat-2#Show spanning vlan 56

```
VLAN0056
  Spanning tree enabled protocol ieee
  Root ID    Priority    24632
             Address     0011.bbeb.8780
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24632  (priority 24576 sys-id-ext 56)
             Address     0011.bbeb.8780
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
```

```
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
```

Task 12

Cat-1 should be configured such that the ports that routers R1 to R6 are connected will bypass listening and learning state. If any of these ports receive BPDU packets, they should transition into errdisable state. Use minimum number of commands to accomplish this task. This configuration should only be applied to the ports that the routers R1 - R6 are connected to as well as any future port that has this feature enabled.

On Cat-1

Cat-1(config)#**Spanning-tree portfast bpduguard default**

Cat-1(config)#Interface range F0/1 - 6
Cat-1(config-if)#Spanning-tree portfast

Once the "Spanning-tree portfast" command is entered you should see the following warning message:

%Warning: portfast should only be enabled on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, etc... to this interface when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION
***%Portfast will be configured in 6 interfaces** due to the range command but will only have effect when the interfaces are in a non-trunking mode.*

The "spanning-tree portfast bpduguard default" command in global config mode will shut the port down in err-disable mode if any portfast enabled port receives BPDU packets.

To verify the configuration:

On Cat-1

Cat-1#Sh spanning-tree interface f0/1 portfast

VLAN0012 enabled

Note if the output of the above show command states "no spanning tree info available for FastEthernet0/1", it only means that the F0/0 interface of R1 is in Shutdown mode.

To test the configuration:

On SW2

Cat-2(config)#**spanning-tree portfast bpduguard default**

Cat-2(config)#**int f0/23**
Cat-2(config-if)#**swi mode acc**
Cat-2(config-if)#**spanning-tree portfast**
Cat-2(config-if)#**No shut**

Note if the f0/23 interface of Switch 3 is enabled, it will generate BPDUs and because of this configuration, F0/23 interface of SW-2 will transition into err-disable mode, as follows:

On Switch 3

Switch(config)#**int f0/23**
Switch(config-if)#**NO shut**

On Cat-2

You should see the following messages:

%SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port FastEthernet0/23 with BPDU Guard enabled. Disabling port.
%PM-4-ERR_DISABLE: bpduguard error detected on Fa0/23, putting Fa0/23 in err-disable state

To verify that interface f0/23 is in err-disable mode:

On Cat-2

Cat-2#Sh inter f0/23 status

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/23		err-disabled	1	auto	auto	10/100BaseTX

To change the configuration back:

On Cat-2

```
Cat-2(config)#NO spanning-tree portfast bpduguard default

Cat-2(config)#int f0/23
Cat-2(config-if)#Shut
Cat-2(config-if)#NO spanning-tree portfast
```

Task 13

Cat-2 should be configured such that the ports that routers R1 to R6 are connected (F0/1 - F0/6) will bypass listening and learning state. If any of these ports receive BPDU packets, they should no longer bypass their listening and learning state. This configuration should apply to existing and future ports that are configured as portfast.

On Cat-2

```
Cat-2(config)#Spanning-tree portfast bpdufilter default

Cat-2(config)#Interface range F0/1 - 6
Cat-2(config-if)#Spanning-tree portfast
```

When BPDUFilter is enabled globally, it will apply to all portfast enabled interfaces; If any portfast enabled interface receives BPDUs, it will bypass listening and learning state, which means that it will loose it's portfast state.

Task 14

You received a request from the IT department to monitor and analyze all the packets sent and received by the host connected to port F0/14 on Cat-1; you have connected the packet analyzer to port F0/15 on the same switch. Configure the switch to accommodate this request.

On Cat-1

Cat-1(config)#**monitor session 1 source interface F0/14 both**
Cat-1(config)#**monitor session 1 destination interface F0/15**

Note the following:

- There can only be two monitor sessions configured on a given switch
- Their direction to monitor can be configured as Rx, Tx, or Both. Rx is for received traffic, Tx is for Transmitted traffic, and both is in both directions. Both is the default direction.
- To verify Enter the "Show monitor session 1" command.

To verify the configuration:

On Cat-1

Cat-1#Show monitor session 1

```
Session 1
---------
Type              : Local Session
Source Ports      :
    Both          : Fa0/14
Destination Ports : Fa0/15
    Encapsulation : Native
          Ingress : Disabled
```

Task 15

You received another request from your IT department to keep track of all the MAC addresses that are learned by Cat-2 port F0/18. The switch must use the NMS located at 192.168.1.1 /24; this switch should send a community string of "Private" with the notification operation. You should use an IP address of 2.2.2.2 /8 to accomplish this task.

On Cat-2

Cat-2(config)#**Snmp-server host 192.168.1.1 traps Private**

%IP_SNMP-3-SOCKET: can't open UDP socket
Unable to open socket on port 161

Note since this switch is not configured with an IP address, it will fail to configure

the Snmp server. Therefore, an IP address should be configured before entering the "snmp-server" command as follows:

Cat-2(config)#Int lo0
Cat-2(config-if)#Ip addr 2.2.2.2 255.0.0.0

The following command identifies he NMS and sends a community string of Private with the notification operation.

Cat-2(config)#**snmp-server host 192.168.1.1 traps** Private

The following command configures the switch to send mac-address traps to the NMS:

Cat-2(config)#**snmp-server enable traps mac-notification**

Cat-2(config)#Inter f0/18
Cat-2(config-if)#**snmp trap mac-notification** added

The above command enables the SNMP trap on interface F0/18 and configures the switch to send MAC notification traps whenever a MAC-address is added. If the switch must be configured to report the MAC addresses that are learnt and expired, then "snmp trap mac-notification change removed" command must also be configured.

To verify the configuration:

On Cat-2

Cat-2#Show mac-address-table notification inter f0/18

MAC Notification Feature is Disabled on the switch

Interface	MAC Added Trap	MAC Removed Trap
FastEthernet0/18	Enabled	Disabled

Note the mac-notification is disabled, the following command will enable the mac-notification on the switch:

Cat-2(config)#**mac address-table notification**

To verify the configuration:

On Cat-2

Cat-2#Show mac-address-table notification interface F0/18

MAC Notification Feature is Enabled on the switch

Interface	MAC Added Trap	MAC Removed Trap
FastEthernet0/18	**Enabled**	Disabled

Task 16

Configure Cat-2's port F0/14 to limit the amount of bandwidth utilization for broadcast traffic to 50%.

On Cat-2

Cat-2(config)#Interface F0/14
Cat-2(config-if)#**Storm-control broadcast level 50.00**

Storm-control can be used for Broadcast, Unicast and Multicast traffic, this command specifies traffic suppression level for a given type of traffic for a particular interface. The level can be from 0 to 100 and an optional fraction of a level can also be configured from 0 – 99. A threshold value of 100 percent means that no limit is placed for the specified type of traffic; a value of 0.0 means that the particular type of traffic is blocked all together.

On 3550 switches when the rate of Multicast traffic exceeds a predefined threshold, all incoming traffic (Broadcast, Multicast and Unicast) is dropped until the level of Multicast traffic is dropped below the threshold level. Once this occurs, only the Spanning-tree packets are forwarded. When Broadcast or Unicast thresholds are exceeded, traffic is blocked for only the type of traffic that exceeded the threshold.

To verify the configuration:

On Cat-2

Cat-2#Show storm-control f0/14 broadcast

Interface	Filter State	Upper	Lower	Current

Fa0/14 Forwarding 50.00% 50.00% 0.00%

If you get "Link Down" as Filter State, the port might be down.

Task 17

Mac addresses learnt dynamically by these two switches should not stay in the MAC address table if they are inactive for longer than 10 minutes.

By default the MAC addresses that are inactive will expire within 300 seconds, this task is asking for a 10 minutes threshold, 10 minutes equates to 600 seconds; the following command sets the idle timer to 10 minutes:

On Both Switches

(config)#**Mac address-table aging-time 600**

To verify the configuration:

On Both Switches

#Sh mac address-table aging-time

```
Vlan    Aging Time
----    ----------
  1       600
 12       600
 34       600
 56       600
```

Task 18

For management purposes, assign an IP address of 10.1.1.11 /24 to Cat-1, with a default gateway of 10.1.1.100 /24.

On Cat-1

```
Cat-1(config)#Inter Vlan 1
Cat-1(config-if)#Ip address 10.1.1.11 255.255.255.0
Cat-1(config-if)#No shut

Cat-1(config)#Ip default-gateway 10.1.1.100
```

To verify the configuration:

On Cat-1

Cat-1#Sh ip interface vlan 1

```
Vlan1 is up, line protocol is up
  Internet address is 10.1.1.11/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
```

(The rest of the output is omitted)

Cat-1#Sh ip route

```
Default gateway is 10.1.1.100

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

Task 19

Configure routers R1 and R3 using the following IP addresses:

- R1 - F0/0 = 10.1.12.1 /24
- R3 - F0/0 = 10.1.34.3 /24

Configure Cat-1 to route between VLAN 12 and 34, use ping to verify the communication. The gateway for VLAN12 should be configured to be 10.1.12.100, and the gateway for VLAN 34 should be configured to be 10.1.34.100.

On R1

```
R1(config)#Interface F0/0
R1(config-if)#Ip address 10.1.12.1 255.255.255.0
```

R1(config-if)#No shut

R1(config)#Ip route 0.0.0.0 0.0.0.0 10.1.12.100

On R3

R3(config)#Interface F0/0
R3(config-if)#Ip address 10.1.34.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Ip route 0.0.0.0 0.0.0.0 10.1.34.100

On Cat-1

Cat-1(config)#**Ip routing**

Cat-1(config)#Interface Vlan 12
Cat-1(config-if)#Ip address 10.1.12.100 255.255.255.0

Cat-1(config)#Interface Vlan 34
Cat-1(config-if)#Ip address 10.1.34.100 255.255.255.0

A Switch Virtual Interface (SVI) represents a VLAN of switch ports as one interface to the routing. Only one SVI can be associated with a VLAN. This is necessary when configuring InterVlan routing.
When creating an SVI for a VLAN, the designated number must match the VLAN number.

To verify the configuration:

On R1

R1#Ping 10.1.34.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

On R3

R3#Ping 10.1.12.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Note By default, IP routing is disabled on the switch and if the "IP Routing" command is NOT enabled on Cat-1, the communication between R1 and R3 can NOT occur.

Task 20

Remove the configuration from the previous step and configure InterVlan routing between VLANs 12 and 34. DO NOT use SVIs to accomplish this task. F0/1 interface of any router can be used to accomplish this task. Use the IP addressing from the previous task. Ensure to use an industry standard protocol/s to accomplish this task.

Since R5's F0/0 is part of VLAN 56, R5's F0/1 is used to accomplish this task.

On Cat-1

```
Cat-1(config)#NO Interface Vlan 12
Cat-1(config)#NO Interface Vlan 34
```

On Cat-2

```
Cat-2(config)#Interface F0/5
Cat-2(config-if)#Switchport trunk encap Dot1q
Cat-2(config-if)#Switchport mode trunk
```

On R5

```
R5(config)#Interface F0/1
R5(config-if)#No shut

R5(config)#Int f0/1.12
R5(config-if)#Encap dot1q 12
R5(config-if)#Ip address 10.1.12.100 255.255.255.0

R5(config)#Int f0/1.34
R5(config-if)#Encap dot1q 34
R5(config-if)#Ip address 10.1.34.100 255.255.255.0
```

To verify the configuration:

On R1 & R3

R1#Clear arp

On R1

R1#Ping 10.1.34.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

On R3

R3#Ping 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Task 21

Configure Cat-1 such that whenever the switch learns or removes a MAC address on its port F0/4, an SNMP notification is generated and sent to the NMS located at 192.168.1.1 with a community string of CAT1. Since there are many users coming and going from the network, set up a trap interval time to bundle the notification traps and reduce network traffic using the following parameters:

- The traps should be generated every 30 minutes.
- The trap should contain a maximum of 150 entries.

This feature enables us to track users on a network by storing the Mac address activity on the switch. Once configured, every time a MAC address is learned or removed an SNMP notification is generated and sent to the NMS. On a very busy network when lots of users come and go, the default behavior is that an SNMP trap is sent every second. Because this can consume bandwidth, there

are two parameters that can be configured to remedy this situation and they are as follows:

- **Mac address-table notification interval** – This value specifies the notification trap interval in seconds between each set of traps that are generated to the NMS. Default value is one second, and the range is 0 – 2,147,483,647 seconds.

- **Mac address-table notification history-size** – Specifies the maximum number of entries in the MAC notification history table. The default value is 1, and the range is 1 – 500 entries.

On Cat-1

```
Cat-1(config)#Snmp-server host 192.168.1.1 traps CAT1
Cat-1(config)#Snmp-server enable traps mac-notification
Cat-1(config)#Mac-address-table notification
Cat-1(config)#Mac-address-table notification interval 1800
Cat-1(config)#Mac-address-table notification history-size 150

Cat-1(config)#Int f0/4
Cat-1(config-if)#Snmp trap mac-notification added
Cat-1(config-if)#Snmp trap mac-notification removed
```

To verify the configuration:

On Cat-1

```
Cat-1#Show mac-address-table notification interface f0/4

MAC Notification Feature is Enabled on the switch

Interface          MAC Added Trap       MAC Removed Trap
---------          ------------------   ----------------------
FastEthernet0/18   Enabled              Enabled

Cat-1#Show mac-address-table notification

MAC Notification Feature is Enabled on the switch

Interval between Notification Traps : 1800 secs
Number of MAC Addresses Added : 0
Number of MAC Addresses Removed : 0
Number of Notifications sent to NMS : 0
```

Maximum Number of entries configured in History Table : 150
Current History Table Length : 0
MAC Notification Traps are Enabled
History Table contents

To verify the configuration:

On R4

R4(config)#int f0/0
R4(config-if)#IP address 4.4.4.4 255.0.0.0
R4(config-if)# no shut

R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note the purpose of the above configuration is to generate some traffic. The following Show command reveals that one MAC address was learned and added to the table.

On SW1

Cat-1#Sh mac-address-table notification

MAC Notification Feature is Enabled on the switch
Interval between Notification Traps : 1800 secs
Number of MAC Addresses Added : 1
Number of MAC Addresses Removed : 0
Number of Notifications sent to NMS : 0
Maximum Number of entries configured in History Table : 150
Current History Table Length : 0
MAC Notification Traps are Enabled
History Table contents

On R4

R4(config)#int f0/0
R4(config-if)#Shut

The output of the following show command reveals that one MAC address was removed.

On Cat-1

Cat-1#Sh mac-address-table notification

MAC Notification Feature is Enabled on the switch
Interval between Notification Traps : 1800 secs
Number of MAC Addresses Added : 1
Number of MAC Addresses Removed : 1
Number of Notifications sent to NMS : 0
Maximum Number of entries configured in History Table : 150
Current History Table Length : 0
MAC Notification Traps are Enabled
History Table contents

Task 22

Optimize Cat-1 using the following policies:
Cat-1 should be configured such that its memory resources are optimized for routing.

Switch database management (SDM) are templates that can be configured to allocate memory resources in the switch for a specific feature depending on what the switch is used for in a given network.
A switch can be configured to use one of the following templates:

- **Access** – Used for QOS classification and Security.
- **Routing** – Used for routing
- **Vlan** – Disables routing and sets the switch to be a layer 2 switch.
- **Extended-match** – reformats routing memory space to allow 144-bit layer 3 TCAM support needed for WCCP and/or multiple VRF instances.

On Cat-1

Cat-1(config)#Sdm prefer routing

You must reboot for these settings to take effect.

Cat-1#WR

Cat-1#Reload

To Verify the configuration after the reload:

On Cat-1

Cat-1#Show sdm prefer

```
The current template is "desktop routing" template.
The selected template optimizes the resources in
the switch to support this level of features for
8 routed interfaces and 1024 VLANs.
 number of unicast mac addresses:                  3K
 number of IPv4 IGMP groups + multicast routes:    1K
 number of IPv4 unicast routes:                    11K
 number of directly-connected IPv4 hosts:          3K
 number of indirect IPv4 routes:                   8K
 number of IPv4 policy based routing aces:         512
 number of IPv4/MAC qos aces:                      512
 number of IPv4/MAC security aces:                 1K
```

On Cat-2

Cat-2#Sh sdm prefer

```
The current template is "desktop default" template.
The selected template optimizes the resources in
the switch to support this level of features for
8 routed interfaces and 1024 VLANs.
  number of unicast mac addresses:                  6K
  number of IPv4 IGMP groups + multicast routes:    1K
  number of IPv4 unicast routes:                    8K
  number of directly-connected IPv4 hosts:          6K
  number of indirect IPv4 routes:                   2K
  number of IPv4 policy based routing aces:         0
  number of IPv4/MAC qos aces:                      512
  number of IPv4/MAC security aces:                 1K
```

Note, the difference in memory allocation is revealed if the buffer allocation of Cat-2 is compared to the Cat-1.

Task 23

Create VLANs 30, 31 and 32 on Cat-1 and ensure that these VLANs can not traverse the trunk link between Cat-1 and Cat-2.

By default a trunk port sends and receives traffic from all VLANs, however, a given VLAN or VLANs can be removed from the trunk link in order to prevent traffic from that VLAN/s from traversing over the trunk.

On Cat-1

Cat-1(config)#Vlan 30-32
Cat-1(config-vlan)#exit

Before configuring the task we have to check to see if the VLANs that we just created can traverse the trunk link.

Cat-1#Show interface trunk

```
Port      Mode        Encapsulation  Status       Native vlan
Fa0/19    on          isl            trunking     1
Fa0/20    on          isl            trunking     1

Port      Vlans allowed on trunk
Fa0/19     1-4094
Fa0/20     1-4094

Port      Vlans allowed and active in management domain
Fa0/19    1,12,30-32,34,56

Fa0/20    1,12,30-32,34,56

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/19     1,12,34,56
Fa0/20     1
```

To remove those VLANs from the trunk links:

On Both Switches

(config)#Interface range f0/19-20
(config-if-range)#**Switchport trunk allowed vlan except 30,31,32**

Note if an EtherChannel was created, the command had to be configured directly

under the port-channel interface.

To Verify the configuration:

On Cat-1

Cat-1#Show int trunk

```
Port      Mode      Encapsulation  Status      Native vlan
Fa0/19     on         isl          trunking     1
Fa0/20     on         isl          trunking     1
Port      Vlans allowed on trunk

Fa0/19     1-29,33-4094         Note VLANs 30 – 32 are removed from the trunk
Fa0/20     1-29,33-4094

Port      Vlans allowed and active in management domain
Fa0/19     1,12,34,56
Fa0/20     1,12,34,56

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/19     1,12,34,56
Fa0/20     1
```

Note the options that can be used with "Switchport trunk allowed VLAN" command are: Remove, add, all, and except.
The "Switchport trunk allowed vlan remove 30,31,32" command could accomplish the same task.

Task 24

Configure Cat-1's port F0/15 and F0/16 such that when client PCs connect to these ports, they automatically become member of a given VLAN. Cat-1 should be configured to use 10.1.1.1 as the primary and 10.1.1.2 as the secondary VMPS server. Ensure that the local switch reconfirms the VLAN membership every half hour and if the VMPS can not be contacted, the local switch will retry 5 times before considering the VMPS unavailable.

VMPS:

- **The 3550 switch can't be setup as a VMPS server, but it can be configured as a VMPS client.**

➢ The client communicates with the VMPS through Vlan Query Protocol (VQP).

➢ When a VMPS receives a VQP from the client, it searches its database for a MAC to VLAN mapping, and if the mapping is found, it conveys the VLAN information to the client and then the client assigns that given VLAN to the port that the client is connected to.

➢ The VMPS can operate in Secure mode, which means that if a MAC to VLAN mapping can not be found in its database, the VMPS will send a port-shutdown-message to the client and the client will shut down that given port, however, if the VMPS is not configured in a secure mode, it will send access-deny message, and the client will constantly monitor the port and will reject all traffic from that given port.

➢ The VMPS client periodically reconfirms the VLAN membership information received from the VMPS server. By default this is performed every 60 minutes, this interval can be changed using "VMPS reconfirm" global config command.

➢ If the VMPS client can't contact the VMPS server, it will retry to establish that communication three times and this value can be changed using vmps retry" command in the global config mode.

➢ The database is in form of an ASCII file saved on a TFTP server, which the VMPS server accesses.

On Cat-1

Before configuring this task we should check some of the default values:

```
Cat-1#Show vmps

VQP Client Status:
--------------------
VMPS VQP Version:   1
Reconfirm Interval: 60 min

Server Retry Count: 3
VMPS domain server:
Reconfirmation status
---------------------
VMPS Action:        No Dynamic Port
```

VMPS VQP version is version 1, and the reconfirmation is at its default value of 60 minutes, and the retry value is set to 3. There are no VMPS servers.

Cat-1(config)#int range f0/15 - 16
Cat-1(config-if-range)#switchport mode access
Cat-1(config-if-range)#switchport access vlan dynamic
Cat-1(config-if-range)#no shut

The above command sets ports F0/15 and F0/16 to VLAN dynamic, which means that they will acquire their VLAN information dynamically. The "no shut" command is required because these ports were shut down earlier

Cat-1(config)#vmps reconfirm 30
Cat-1(config)#vmps retry 5

The above two commands configure the reconfirmation interval to 30 minutes and the retry counter to 5.

Cat-1(config)#vmps server 10.1.1.1 primary
Cat-1(config)#vmps server 10.1.1.2

These commands configure the primary and the secondary VMPS servers.

To verify the configuration:

On Cat-1

Cat-1#Show vmps

```
VQP Client Status:
--------------------
VMPS VQP Version:   1
Reconfirm Interval:     30 min
Server Retry Count:     5
VMPS domain server: 10.1.1.2
                    10.1.1.1 (primary, current)
Reconfirmation status
---------------------------------------------------------------------
VMPS Action:        No Dynamic Port
```

Task 25

Port F0/17 on Cat-1 is connected to a Cisco 7960 IP Phone. Voice traffic that originates from the phone is tagged with a CoS of 5.

A PC is connected to the 7960 IP Phone which is generating traffic with CoS of 3. Ensure that the data traffic belongs to VLAN 3 and the Voice traffic belongs to VLAN 5. The traffic originated by the 7960 IP Phone should maintain it's CoS value, whereas, the traffic that originated from the PC connected to the 7960 IP Phone should be re-written with a CoS of 1.

On Cat-1

Cat-1(config)#Mls qos

Cat-1(config)#Interface F0/17
Cat-1(config-if)#Switchport access Vlan 3
Cat-1(config-if)#Switchport voice Vlan 5
Cat-1(config-if)#Switchport priority extend cos 1
Cat-1(config-if)#Mls qos trust cos
Cat-1(config-if)#No shut

When the phone gets connected to the switch it will form an 802.1q trunk link. The traffic destined to the PC will be carried in the access VLAN, whereas the traffic destined for the 7960 IP Phone will be carried in Voice VLAN.
By default the 3550 doesn't process the CoS value and rewrites all frames with a CoS value of 0. To configure the phone such that it processes the CoS values, the QOS must be enabled globally using the "mls qos" command.
To configure the switch so it trusts the incoming CoS value from the 7960 IP Phone the "mls qos trust cos" command is used.
Since the PC connected to the IP Phone can send traffic to the Phone with any Cos value and the phone wants to ensure that the voice traffic that it generates get better priority, it overrides the CoS for all traffic that is originated by the

PC. In this task we have to configure the switch such that it re-writes the traffic with a CoS of 1, therefore, the "Switchport priority extended cos 1" command is used. The "no shut" command is required because the port was shut down earlier.

Task 26

Configure trunking between Cat-1 and Cat-2 such that VLAN 12 does not get tagged when the traffic for this VLAN traverses the trunk.

Note the trunking encapsulation on the trunk links should have been DOT1Q; in the CCIE lab, when configuring a given section, the entire section should be read before configuring the individual tasks within that section.

When a trunk is configured with Dot1q, it can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic in the native VLAN ONLY. If a given VLAN should NOT be tagged as it traverses the trunk link then, that VLAN should be set as the native VLAN.
When the native VLAN is changed, ensure that the change is configured on both switches or the trunk link will go down.

On Both Switches

```
(config)#Interface range F0/19-20
(config-if-range)#Switchport trunk encap dot1q
```

To Verify the configuration:

On Cat-1

Cat-1#Show int trunk

```
Port      Mode      Encapsulation  Status     Native vlan
Fa0/19     on        802.1q        trunking     1
Fa0/20     on        802.1q        trunking     1

Port     Vlans allowed on trunk
Fa0/19     1-29,33-4094
Fa0/20     1-29,33-4094


Port      Vlans allowed and active in management domain
Fa0/19     1,3,5,12,34,56
Fa0/20     1,3,5,12,34,56

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/19     1,3,5,12,34,56
Fa0/20     1
```

To configure the native VLAN:

On Both Switches

```
(config)#Interface range F0/19-20
(config-if-range)#Switchport trunk native VLAN 12
```

To verify the configuration:

On Cat-1

Cat-1#Show interface trunk

```
Port      Mode    Encapsulation  Status     Native vlan
Fa0/19    on      802.1q         trunking   12
Fa0/20    on      802.1q         trunking   12

Port      Vlans allowed on trunk
Fa0/19    1-29,33-4094
Fa0/20    1-29,33-4094

Port      Vlans allowed and active in management domain
Fa0/19    1,3,5,12,34,56
Fa0/20    1,3,5,12,34,56

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/19    1,3,5,12,34,56
Fa0/20    1
```

On Cat-2

Cat-2#Show interface trunk

```
Port      Mode    Encapsulation  Status     Native vlan
Fa0/5     on      802.1q         trunking   1
Fa0/19    on      802.1q         trunking   12
Fa0/20    on      802.1q         trunking   12

Port      Vlans allowed on trunk
Fa0/5     1-4094
Fa0/19    1-29,33-4094
Fa0/20    1-29,33-4094

Port      Vlans allowed and active in management domain
Fa0/5     1,3,5,12,30-32,34,56
Fa0/19    1,3,5,12,34,56
Fa0/20    1,3,5,12,34,56

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/5     1,3,5,12,30-32,34,56
Fa0/19    1,12,34,56
Fa0/20    none
```

Task 27

The IT department decided to stop monitoring port F0/14 from Task 14, you have received a new request to monitor port F0/14 on Cat-1 but the protocol analyzer is connected to port F0/18 on Cat-2. Configure the switches to accommodate this request.

On Cat-1

```
Cat-1(config)#NO monitor session 1

Cat-1(config)#Vlan 90
Cat-1(config-vlan)#Remote-span
Cat-1(config-vlan)#Exit
```

The creation of this VLAN can only be done in the global configuration mode, because this is the only mode that allows us to set the VLAN as remote-span. Ensure that this VLAN is propagated to Cat-2.

To verify the configuration:

On Cat-1:

Cat-1#Sh vlan brie

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12, Fa0/13, Fa0/14
                                                Fa0/18, Fa0/21, Fa0/22, Fa0/23
                                                Fa0/24, Gi0/1, Gi0/2
3    VLAN0003                         active    Fa0/17
5    VLAN0005                         active    Fa0/17
12   VLAN0012                         active    Fa0/1, Fa0/2
30   VLAN0030                         active
31   VLAN0031                         active
32   VLAN0032                         active
34   VLAN0034                         active    Fa0/3, Fa0/4
56   VLAN0056                         active    Fa0/5, Fa0/6
90   VLAN0090                         active  <------  Ensure that this VLAN is propagated
                                                       to Cat-2
```

(The rest of the output is omitted)

On Cat-2

Cat-2#Sh vlan brie

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4

                                                Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                                Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                                Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                                Fa0/18, Fa0/21, Fa0/22, Fa0/23
                                                Fa0/24, Gi0/1, Gi0/2
3    VLAN0003                         active
5    VLAN0005                         active
12   VLAN0012                         active
30   VLAN0030                         active
31   VLAN0031                         active
32   VLAN0032                         active
34   VLAN0034                         active
56   VLAN0056                         active
90   VLAN0090                         active    <---------    Note the VLAN is propagated.
```

(The rest of the output is omitted)

On Cat-1

Cat-1#Show vlan remote-span

```
Remote SPAN VLANs
------------------------------
90
```

On Cat-2

Cat-2#Show vlan remote-span

```
Remote SPAN VLANs
------------------------------
90
```

Note VLAN 90 should be displayed as remote-span on both switches.

On Cat-1

Cat-1(config)#**Monitor session 1 source interface F0/14**

Cat-1(config)#**Monitor session 1 destination remote vlan 90**

To verify the configuration:

On Cat-1

Cat-1#Show monitor session 1

```
Session 1
---------
Type                    : Remote Source Session
Source Ports            :
    Both                : Fa0/14
Dest RSPAN VLAN         : 90
```

On Cat-2

Cat-2(config)#**Monitor session 1 source remote vlan 90**
Cat-2(config)#**Monitor session 1 destination interface F0/18**

Port F0/18 is where the protocol analyzer is connected.

To verify the configuration:

On Cat-2

Cat-2#Sh monitor session 1

```
Session 1
---------
Type                    : Remote Destination Session
Source RSPAN VLAN       : 90
Destination Ports       : Fa0/18
   Encapsulation        : Native
       Ingress          : Disabled
```

RSPAN extends SPAN by enabling remote monitoring of multiple switches across your network. The traffic for RSPAN traverses over a user defined RSPAN VLAN (remote vlan), in this case VLAN 90. The SPAN traffic from port F0/14 is reflected to VLAN 90 (The RSPAN VLAN) and then forwarded over the trunk to port F0/18 an RSPAN destination.

Task 28

Configure the hostname of the third switch to be Cat-3, and disable all ports but F0/21-22. This Switch should join the "CCIE" VTP domain.

On the third Switch

Switch(config)#**Hostname Cat-3**

Cat-3(config)#int range f0/1 - 20 , F0/23 - 24
Cat-3(config-if-range)#Shut

Cat-3(config)#vtp domain CCIE
Cat-3(config)#vtp password Cisco

Note sometimes a VLAN needs to be created in order to propagate the existing VLANs, as follows:

On Cat-3

Cat-3(config)#vlan 99
Cat-3(config-vlan)#exit

Note the VLANs are propagated:

Cat-3#Sh vlan brie

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active
30   VLAN0030                         active
31   VLAN0031                         active
32   VLAN0032                         active
34   VLAN0034                         active
56   VLAN0056                         active
90   VLAN0090                         active
```

Next, Vlan 99 is removed:

```
Cat-3(config)#NO vlan 99

Cat-3#Show vlan bric | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active
30   VLAN0030                         active
31   VLAN0031                         active
32   VLAN0032                         active
34   VLAN0034                         active
56   VLAN0056                         active
```

Task 29

Configure ports F0/21 and F0/22 on Cat-3 and Cat-1 as trunk links using an industry standard protocol, these links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should NOT negotiate by using any protocol to accomplish this task.

EtherChannels provide the follows:

- Fault-tolerant, high speed links between switches and routers.
- EtherChannel provides an automatic recovery for the loss of a link by redistributing the traffic across the remaining link/s.
- STP will not block one of the links in the bundle because to STP, the bundle looks like a single link.
- Up to 8 links can be combined to provide more bandwidth.
- The links within the bundle must have the same characteristics such as duplexing, speed and etc.
- EtherChannel can be configured as layer 2 or layer 3.
- With Layer 3, a logical interface (Port-Channel) is statically configured and all Layer 3 configurations are performed under that interface.
- With Layer 2, the logical interface is created automatically.
- With both Layer 2 and Layer 3, physical interfaces must be manually assigned to

the logical interface using "channel-group" configuration command.

- EtherChannels can be configured automatically using Port aggregation protocol (PAgP) or Link Aggregation protocol (LACP).
- PAgP is a Cisco proprietary protocol, whereas LACP is an industry standard IEEE 802.3ad protocol.
- Switches can be configured to use PAgP by configuring them in AUTO or DESIRABLE mode.
- Switches can be configured to use LACP by configuring them in ACTIVE or PASSIVE mode.
- If the switches are configured in ON mode, they will not exchange LACP or PAgP packets.

There are 5 modes that the switches can be configured in:

- ON – Forces the interface into an EtherChannel without PAgP or LACP packets, both switches must be configured in ON mode for the EtherChannel to be established.

- ACTIVE – Used in LACP, the switches will actively negotiate an EtherChannel link.

- PASSIVE – Used in LACP, it places the interface in a passive negotiation mode where it only responds to LACP packets that it receives. In this mode the switch will not start the negotiation process; this setting minimizes the transmission of LACP packets.

- AUTO – Used in PAgP, it places the interface in a passive negotiation mode; It only responds to PAgP packets that it receives. In this mode the switch will not start the negotiation process; this setting minimizes the transmission of PAgP packets.

- DESIRABLE – Used in PAgP, the switches will actively negotiate an EtherChannel link.

The following table is very important when configuring EtherChannels:

Switch one is configured as	Switch two is configured as	Will an EtherChannel be established?
Desirable	Desirable	YES
Desirable	Auto	YES
Auto	Auto	NO
Active	Active	YES
Active	Passive	YES
Passive	Passive	NO

Before configuring EtherChannel, you should check to ensure that the interfaces are configured with the same characteristics.
The best way to configure an EtherChannel is to configure the Channel-group under the interfaces first, as follows:

On Both Switches

```
(config)#Int range f0/21 - 22
 (config-if-range)#Channel-group 1 mode on
(config-if-range)#no shut
```

Then, configure the port-channel that is created automatically as trunk.

```
(config-if-range)#int port-channel 1
(config-if)#switchport trunk encapsulation dot1q
(config-if)#Switchport mode trunk
```

To verify the configuration:

On Cat-1

```
Cat-1#Show int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      12
Fa0/20      on           802.1q         trunking      12
Po1         on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-29,33-4094
Fa0/20      1-29,33-4094
Po1         1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,12,34,56,90
Fa0/20      1,12,34,56,90
Po1         1,12,30-32,34,56,90

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
Fa0/20      1,12,34,56,90
Po1         1,12,30-32,34,56,90
```

On Cat-3

Cat-3#Sh int trunk

```
Port      Mode       Encapsulation  Status       Native vlan
Po1       on         802.1q         trunking     1

Port      Vlans allowed on trunk
Po1       1-4094

Port      Vlans allowed and active in management domain
Po1       1,12,30-32,34,56

Port      Vlans in spanning tree forwarding state and not pruned
Po1       1,12,30-32,34,56
```

On Cat-1

Cat-1#Show spanning-tree int f0/21

```
Mst Instance      Role  Sts     Cost     Prio.Nbr Type
---------------- -------- -------- ------- --------------------------------
VLAN0001         Root FWD 12       128.616  P2p
VLAN0012         Root FWD 12       128.616  P2p
VLAN0030         Root FWD 12       128.616  P2p
VLAN0031         Root FWD 12       128.616  P2p
VLAN0032         Root FWD 12       128.616  P2p
VLAN0034         Root FWD 12       128.616  P2p
VLAN0056         Root FWD 12       128.616  P2p
```

Cat-1#Show spanning-tree int f0/22

```
Mst Instance      Role  Sts     Cost     Prio.Nbr Type
---------------- -------- -------- ------- --------------------------------
VLAN0001         Root FWD 12       128.616  P2p
VLAN0012         Root FWD 12       128.616  P2p
VLAN0030         Root FWD 12       128.616  P2p
VLAN0031         Root FWD 12       128.616  P2p
VLAN0032         Root FWD 12       128.616  P2p
VLAN0034         Root FWD 12       128.616  P2p
VLAN0056         Root FWD 12       128.616  P2p
```

On Cat-3

Cat-3#Show spanning-tree int f0/21

```
Vlan             Role  Sts   Cost     Prio.Nbr Type
---------------- ----  ----- -------- -------- ----------------
VLAN0001         Desg FWD 12       128.65   P2p
VLAN0012         Desg FWD 12       128.65   P2p
VLAN0030         Desg FWD 12       128.65   P2p
VLAN0031         Desg FWD 12       128.65   P2p
VLAN0032         Desg FWD 12       128.65   P2p
VLAN0034         Desg FWD 12       128.65   P2p
VLAN0056         Desg FWD 12       128.65   P2p
```

Cat-3#Show spanning-tree int f0/22

```
Vlan             Role  Sts   Cost     Prio.Nbr Type
---------------- ----  ----- -------- -------- ----------------
VLAN0001         Desg FWD 12       128.65   P2p
VLAN0012         Desg FWD 12       128.65   P2p
VLAN0030         Desg FWD 12       128.65   P2p
VLAN0031         Desg FWD 12       128.65   P2p

VLAN0032         Desg FWD 12       128.65   P2p
VLAN0034         Desg FWD 12       128.65   P2p
VLAN0056         Desg FWD 12       128.65   P2p
```

Note all interfaces are in forwarding state because to spanning-tree the port-channel appears as a single interface.
A "show etherchannel 1 detail" command can reveal that the interfaces are working in the bundle.

Task 30

Ensure that the EtherChannel created in the previous step uses destination MAC addresses to load-balance the traffic load.

Load balancing can be done based on the following:

Source MAC address – Packets forwarded to an EtherChannel are distributed across the ports in the channel based on the source MAC address of the incoming packets. When source MAC address load balancing is enabled, the load distribution based on the source and destination IP address is also enabled

Destination MAC address – If the EtherChannel is between a router and a switch and

since the router has a single MAC address, destination based load balancing is the best way.

To see the default load balancing:

On Cat-1

Note the default load balancing is based on the Source Mac address

Cat-1#show etherchannel load
EtherChannel Load-Balancing Operational State (src-mac):
Non-IP: **Source MAC address**
IPv4: **Source MAC address**
IPv6: **Source IP address**

To configure the load balancing based on the destination Mac addresses:

On Both Switches

(config)#port-channel load-balance dst-mac

To verify the configuration:

Cat-1#show etherchannel load

EtherChannel Load-Balancing Operational State (dst-mac):
Non-IP: Destination MAC address
IPv4: **Destination MAC address**
IPv6: Destination IP address

Task 31

Erase the startup configuration and vlan.dat before proceeding to the next lab

Lab 4
3560 configuration

Task 1

Configure the switches using the following hostnames:
The first switch as Cat-1, the second switch as Cat-2, the third switch as Cat-3 and the forth switch as Cat-4

On the first switch:

```
Switch(config)#ho Cat-1
Cat-1(config)#
```

On the second switch:

```
Switch(config)#ho Cat-2
Cat-2(config)#
```

On the third switch:

```
Switch(config)#ho Cat-3
Cat-3(config)#
```

On the forth switch:

```
Switch(config)#ho Cat-4
Cat-4(config)#
```

Task 2

Configure Cat-1 such that the console messages are displayed with sequence numbers.

On Cat-1

Note to generate a console message all we need to do is go to the global config mode and get back to privilege mode as follows:

Cat-1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Cat-1(config)#end
Cat-1#
00:17:05: %SYS-5-CONFIG_I: Configured from console by console
Note the above console message did not have the sequence numbers, to enable the sequence numbers:

On Cat-1

Cat-1(config)#service sequence-numbers
Cat-1(config)#end
Cat-1#
000057: 00:18:46: %SYS-5-CONFIG_I: Configured from console by console

Note 000057 is the sequence number

Task 3

Disable the timestamps for all console messages including the debug messages on Cat-1

On Cat-1

Cat-1(config)#**NO service timestamps debug**

The above command disables log time stamps, which enables time stamps on log messages showing the time since the system was reloaded for all levels (This is because debug is the default value, so it displays level 7 and all the lower numbers below level 7).

Cat-1(config)#**NO service timestamps log**

The above command disables log time stamps which enables time stamps on log messages showing the time since the system was reloaded.

Cat-1(config)#end
Cat-1#

000058: %SYS-5-CONFIG_I: Configured from console by console

Note there are no time stamps on the above message.

Task 4

Set the time and date of Cat-1 to 16 minutes passed 4 PM, December 26, 2007. The time zone should be set based on Sydney Australia (EST +11). You should use a privilege level and a global config level command to accomplish this task.

On Cat-1

Cat-1#**Clock set 16:16:00 Dec 26 2007**

Cat-1(config)#**clock timezone EST +11**

To verify:

Cat-1#Show clock

16:17:31.972 EST Wed Dec 26 2007

Task 5

Configure Cat-1 such that the system messages are displayed with sequence numbers and current time and date.

On Cat-1

Cat-1(config)#service timestamps log datetime
Cat-1(config)#end

000071: Dec 26 05:19:34: %SYS-5-CONFIG_I: Configured from console by console

Note the sequence number of 000071: followed by the current date and time (Dec 26 05:19:34) is displayed.

Task 6

Configure Cat-1 such that the system messages are displayed with sequence numbers, current date and time in HH:MM:SS and msec and local time and the current timezone.

On Cat-1

Cat-1(config)#**service timestamps log datetime msec localtime show-timezone**
Cat-1(config)#end

000077: Dec 26 16:28:24.354 EST: %SYS-5-CONFIG_I: Configured from console by console

Task 7

Configure Cat-2 using the following policy:

- The switch should log all Emergency, Alerts, Critical, Errors and Warning messages
- The syslog server located at 10.1.1.100.
- The messages should be logged to local4 facility

On Cat-2

Cat-2(config)#logging 10.1.1.100
Cat-2(config)#logging trap 4
Cat-2(config)#logging facility local4

To verify the configuration:

On Cat-2

```
Cat-2#Show logging
Syslog logging: enabled (0 messages dropped, 1 messages rate-limited, 0 flushes, 0 overruns, xml disabled, filtering disable
d)
    Console logging: level debugging, 41 messages logged, xml disabled,
                     filtering disabled
    Monitor logging: level debugging, 0 messages logged, xml disabled,
                     filtering disabled
```

```
Buffer logging: level debugging, 41 messages logged, xml disabled,
            filtering disabled
Exception Logging: size (4096 bytes)
Count and timestamp logging messages: disabled
File logging: disabled
Trap logging: level warnings, 43 message lines logged
   Logging to 10.1.1.100, 0 message lines logged, xml disabled,
        filtering disabled
```

Task 8

Configure Cat-3 to log the system messages to a file called "syslog", this file should be saved in the flash with a max size of 8192. The severity type should be set to "debugging".

On Cat-3

Cat-3(config)#**logging file flash:syslog 8192 debugging**

Cat-3(config)#int f0/1
Cat-3(config-if)#shut
Cat-3(config-if)#NO shut

To verify the configuration:

On Cat-3

```
Cat-3#dir
Directory of flash:/

  2  -rwx        327   Mar 1 1993 00:05:28 +00:00  system_env_vars
  3  -rwx       3426   Mar 1 1993 02:23:17 +00:00  config.you
  4  -rwx       3345   Mar 1 1993 01:49:34 +00:00  config.old
  5  -rwx    7134015   Mar 1 1993 00:04:51 +00:00  c3550-ipservicesk9-mz.122-25.SEE2.bin
  6  -rwx        327   Mar 1 1993 01:25:32 +00:00  syslog
  7  drwx        192   Mar 1 1993 00:03:42 +00:00  c3550-i9q3l2-mz.121-13.EA1a
 24  -rwx          0   Mar 1 1993 00:05:28 +00:00  env_vars
```

Task 9

Configure Cat-1 to disable logging of POE events for it's F0/5 interface..

On Cat-1

```
Cat-1(config)#int F0/5
Cat-1(config-if)#no logging event power-inline-status
```

This command may not be available on your switch if the switch that you are working on is NOT POE (Power Over Ethernet).

Task 10

Configure the system resources of Cat-4 such that unicast routing is disabled and it supports maximum number of Unicast MAC addresses.

On Cat-4

```
Cat-4#Show sdm prefer

 The current template is the default template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1K VLANs.
 number of unicast mac addresses:    5K
 number of igmp groups:              1K
 number of qos aces:                 1K
 number of security aces:            1K
 number of unicast routes:           8K
 number of multicast routes:         1K
```

To change the SDM template for Unicast routing:

```
Cat-4(config)#sdm prefer vlan
```

To verify the configuration:

On Cat-4

```
Cat-4#Show sdm prefer

 The current template is the default template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1K VLANs.
 number of unicast mac addresses:   5K
 number of igmp groups:             1K
 number of qos aces:                1K
 number of security aces:           1K
 number of unicast routes:          8K
 number of multicast routes:        1K

 The template stored for use after the next reload
 is the vlan template.
 The selected template optimizes the resources in

the switch to support this level of features for
 8 routed interfaces and 1K VLANs.
 number of unicast mac addresses:   8K
 number of igmp groups:             1K
 number of qos aces:                1K
 number of security aces:           1K
 number of unicast routes:          0
 number of multicast routes:        0
```

This template disables routing and supports maximum number of Unicast MAC addresses. Typically used for layer 2 switches, if this option is used, routing is done in the software and it severely impacts the switches performance.

Task 11

Configure port F0/1 of Cat-1 as a layer 3 interface and assign an IPv6 address of 12:1:1:12::1/64 to this interface.

On Cat-1

```
Cat-1(config)#int f0/1
Cat-1(config-if)#no switchport
Cat-1(config-if)#ipv6 address 12:1:1:12::1/64
                    ^
```

% Invalid input detected at '^' marker.

Note IPv6 is not enabled and therefore, IPv6 addressing can NOT be assigned to any of the interfaces on this switch, the 3560 switches support IPv6 but the SDM needs to be changed for "dual-ipv4-and-ipv6" before the IPv6 support is enabled.

Cat-1(config)#**sdm prefer dual-ipv4-and-ipv6 default**

Cat-1(config)#int f0/1
Cat-1(config-if)#ipv6 address 12:1:1:12::1/64
Cat-1(config-if)#no shut

To verify the configuration:

Cat-1#sh ipv6 inter f0/1

```
FastEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::217:E0FF:FE26:3B41
  Global unicast address(es):
    12:1:1:12::1, subnet is 12:1:1:12::/64
```

Task 12

Configure F0/23 interface of Cat-1 such that it can detect unidirectional links due to one way traffic on twisted pair. This switch should be configured such that if F0/23 interface transitions into errdisable state, it should automatically recover every 2 minutes and if the port detects unidirectional links it should repeat the cycle again.

On Cat-1 & Cat-4

Cat-x(config)#int f0/23
Cat-x(config-if)#**udld port aggressive**

To verify the configuration:

On Cat-1

Cat-1#Sh udld f0/23

Interface Fa0/23

Port enable administrative configuration setting: Enabled / in aggressive mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Bidirectional
Current operational state: Advertisement - **Single neighbor detected**
Message interval: 15
Time out interval: 5

Entry 1

Expiration time: 43
Device ID: 1
Current neighbor state: Bidirectional
Device name: CHK0649W0TP
Port ID: Fa0/23
Neighbor echo 1 device: FDO1221Z2QT
Neighbor echo 1 port: Fa0/23
Message interval: 15
Time out interval: 5
CDP Device name: SW4

Note if the unidirectional link is detected, the following will be the output of "Show udld F0/23" command:

Interface Fa0/23

Port enable administrative configuration setting: Enabled / in aggressive mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Unknown
Current operational state: Advertisement
Message interval: 7
Time out interval: 5
No neighbor cache information stored

To configure the auto recovery upon detection of unidirectional link:

Cat-1(config)#**errdisable recovery cause udld**

Command enables the timer to automatically recover from the UDLD error-disabled state

Cat-1(config)#**errdisable recovery interval 120**

Command specifies the time to recover from the UDLD error-disabled state

To verify the configuration:

On Cat-1

Cat-1#Sh errdisable recovery

```
ErrDisable Reason    Timer Status
-----------------    ------------
udld                 Enabled
bpduguard            Disabled
security-violatio    Disabled
channel-misconfig    Disabled
vmps                 Disabled
pagp-flap            Disabled
dtp-flap             Disabled
link-flap            Disabled
l2ptguard            Disabled
psecure-violation    Disabled
gbic-invalid         Disabled
dhcp-rate-limit      Disabled
unicast-flood        Disabled
storm-control        Disabled
arp-inspection       Disabled
loopback             Disabled

Timer interval: 120 seconds
```

Task 13

Configure the following IP addresses on Cat-1 and R1:
Cat-1's F0/1 interface – 10.1.1.10 /24, Cat-1 should also have a default gateway pointing to R1.
R1's F0/0 interface 10.1.1.1 /24, Lo0 interface 1.1.1.1 /8, Lo1 interface 100.1.1.1 /24

On R1

```
R1(config)#int f0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0
R1(config-if)#no shut
```

```
R1(config-if)#int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0

R1(config-if)#int lo1
R1(config-if)#ip addr 100.1.1.1 255.255.255.0
```

To verify the configuration:

On R1

R1#Show ip int brie

```
Interface          IP-Address     OK? Method Status                 Protocol
FastEthernet0/0    10.1.1.1       YES manual up                     up
FastEthernet0/1    unassigned     YES unset  administratively down  down
Serial0/0/0        unassigned     YES unset  administratively down  down
Serial0/0/1        unassigned     YES unset  administratively down  down
Loopback0          1.1.1.1        YES manual up                     up
Loopback1          100.1.1.1      YES manual up                     up
```

On Cat-1

```
Cat-1(config)#int f0/1
Cat-1(config)#no switchport
Cat-1(config-if)#ip address 10.1.1.10 255.255.255.0
Cat-1(config-if)#no shut

Cat-1(config)#ip route 0.0.0.0 0.0.0.0 10.1.1.1
```

To verify the configuration:

On Cat-1

Cat-1#Ping 10.1.1.1

```
Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1006 ms
```

Cat-1#Ping 100.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1007 ms

Cat-1#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1006 ms
```

Task 14

Configure a Smartport Macro on Cat-1 such that it pings all the interfaces of R1, this macro should be configured such that it can be executed at any time by entering "TST" in the global config mode.

On Cat-1

```
Cat-1(config)#Macro name TST
Enter macro commands one per line. End with the character '@'.
do Ping 10.1.1.1
do Ping 100.1.1.1
do Ping 1.1.1.1
^Z
Cat-1#
```

To test the configuration:

On Cat-1

```
Cat-1(config)#macro global apply TST

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms

Type escape sequence to abort.
```

Sending 5, 100-byte ICMP Echos to **100.1.1.1**, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to **1.1.1.1**, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

To execute the Macro by just entering "TST", requires configuring an alias, as follows:

Cat-1(config)#**alias configure TST macro global apply TST**

To test the configuration:

On Cat-1

Cat-1(config)#**TST**

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

Task 15

Configure the F0/0 interface of R1-R3 in Vlan 2; configure R1 – R3 based on following parameters:

Router	Interface	IP address	MAC address
R1	F0/0	10.1.1.1 /24	0000.1111.1111
R2	F0/0	10.1.1.2 /24	0000.2222.2222
R3	F0/0	10.1.1.3 /24	0000.3333.3333

On R1

```
R1(config)#int f0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0
R1(config-if)#mac-address 0000.1111.1111
R1(config-if)#no shut
```

On R2

```
R2(config)#int f0/0
R2(config-if)#ip addr 10.1.1.2 255.255.255.0
R2(config-if)#mac-address 0000.2222.2222
R2(config-if)#no shut
```

On R3

```
R3(config)#int f0/0
R3(config-if)#ip address 10.1.1.3 255.255.255.0
R3(config-if)#mac-address 0000.3333.3333
R3(config-if)#no shut
```

On Cat-1

```
Cat-1(config)#int f0/1
Cat-1(config-if)#swi

Cat-1(config)#int range f0/1-3
Cat-1(config-if-range)#swi mode acc
Cat-1(config-if-range)#swi acc v 2
Cat-1(config-if-range)#spanning portfast
```

Task 16

Configure IP source guard on Cat-1 such that it filters traffic based on manually configured IP source bindings. If any of the hosts in this VLAN uses the IP address of another router in this VLAN, the switch (Cat-1) should drop that traffic.

On Cat-1

Cat-1(config)#ip dhcp snooping
Cat-1(config)#ip dhcp snooping vlan 2

The above commands enable DHCP snooping Binding for VLAN 2, these must be configured or else the IP source guard will NOT work.

Cat-1(config)#interface range f0/1-3
Cat-1(config-if-range)#ip verify source

The above command enables Source IP Address Filtering; with "IP Verify Source" command configured under the interfaces, the switch does NOT check the Mac addresses that are bound to the IP addresses.

Cat-1(config)#**ip source binding 0000.1111.1111 vlan 2 10.1.1.1 interface F0/1**
Cat-1(config)#**ip source binding 0000.2222.2222 vlan 2 10.1.1.2 interface F0/2**
Cat-1(config)#**ip source binding 0000.3333.3333 vlan 2 10.1.1.3 interface F0/3**

The above commands configure three entries in the IP Source Bindings table.

To verify the configuration:

On Cat-1

Cat-1#Show ip source binding

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
00:00:22:22:22:22	10.1.1.2	infinite	static	2	FastEthernet0/2
00:00:33:33:33:33	10.1.1.3	infinite	static	2	FastEthernet0/3
00:00:11:11:11:11	10.1.1.1	infinite	static	2	FastEthernet0/1

Total number of bindings: 3

Cat-1#Show ip verify source

Interface	Filter-type	Filter-mode	IP-address	Mac-address	Vlan
Fa0/1	ip	active	10.1.1.1		2
Fa0/2	ip	active	10.1.1.2		2
Fa0/3	ip	active	10.1.1.3		2

To test the configuration:

On R1

R1(config)#int f0/0
R1(config-if)#ip addr 10.1.1.4 255.255.255.0

R1#Ping 10.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note when IP source guard is enabled with source IP address Filtering, IP traffic is filtered based on the source IP address. The Switch forwards IP traffic when the source IP address of that traffic matches an entry in the DHCP snooping binding database or a manually created source binding table.

R1(config)#int f0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0

R1#Ping 10.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

In the above test, the IP address of R1's F0/0 interface was changed to match the configured binding in the switch, therefore, the traffic was allowed.

Task 17

Configure the F0/1 interface of R1-R3 in Vlan 22; configure R1 – R3 based on following parameters:

Router	Interface	IP address	MAC address
R1	F0/1	20.1.1.1 /24	0000.1111.1111
R2	F0/1	20.1.1.2 /24	0000.2222.2222
R3	F0/1	20.1.1.3 /24	0000.3333.3333

On R1

```
R1(config)#int f0/1
R1(config-if)#ip addr 20.1.1.1 255.255.255.0
R1(config-if)#mac-address 0000.1111.1111
R1(config-if)#no shut
```

On R2

```
R2(config)#int f0/1
R2(config-if)#ip addr 20.1.1.2 255.255.255.0
R2(config-if)#mac-address 0000.2222.2222
R2(config-if)#no shut
```

On R3

```
R3(config)#int f0/1
R3(config-if)#ip address 20.1.1.3 255.255.255.0
R3(config-if)#mac-address 0000.3333.3333
R3(config-if)#no shut
```

On Cat-2

```
Cat-2(config)#int range f0/1-3
Cat-2(config-if-range)#swi mode acc
Cat-2(config-if-range)#swi acc v 22
Cat-2(config-if-range)#spanning portfast
```

Task 18

Configure IP source guard on Cat-2 such that it filters traffic based on manually configured IP source and MAC Address Filtering. If the switch detects another MAC or IP address on one of the configured ports, it should drop the traffic.

On Cat-2

```
Cat-2(config)#ip dhcp snooping
Cat-2(config)#ip dhcp snooping vlan 22
```

The above commands enable DHCP snooping Binding for VLAN 22, these must be configured or else the IP source guard will NOT work.

```
Cat-2(config)#interface range f0/1-3
```

Cat-2(config-if-range)#**ip verify source port-security**
Cat-2(config-if-range)#**Switchport port-security**

The above command enables IP Source guard with IP and MAC address filtering. With "IP Verify Source port-security" command configured under the interfaces, the switch will filter based on the MAC and IP addresses. The 'Switchport port-security" command MUST be configured for the interfaces in VLAN 22.

Cat-2(config)#ip source binding 0000.1111.1111 vlan 22 20.1.1.1 interface F0/1
Cat-2(config)#ip source binding 0000.2222.2222 vlan 22 20.1.1.2 interface F0/2
Cat-2(config)#ip source binding 0000.3333.3333 vlan 22 20.1.1.3 interface F0/3

The above commands configure three entries in the IP Source bindings table.

To verify the configuration:

On Cat-2

Cat-2#Show ip source binding

```
MacAddress          IpAddress      Lease(sec)  Type      VLAN   Interface
------------------  -------------  ----------  --------  -----  ----------------
00:00:22:22:22:22   20.1.1.2       infinite    static    22     FastEthernet0/2
00:00:33:33:33:33   20.1.1.3       infinite    static    22     FastEthernet0/3
00:00:11:11:11:11   20.1.1.1       infinite    static    22     FastEthernet0/1
Total number of bindings: 3
```

Cat-2#Show ip verify source

```
Interface  Filter-type  Filter-mode  IP-address     Mac-address         Vlan
---------  -----------  -----------  -------------  ------------------  ---------
Fa0/1      ip-mac       active       20.1.1.1       00:00:11:11:11:11   22
Fa0/2      ip-mac       active       20.1.1.2       00:00:22:22:22:22   22
Fa0/3      ip-mac       active       20.1.1.3       00:00:33:33:33:33   22
```

To test the configuration:

On R1

R1(config)#int f0/1
R1(config-if)#NO mac

R1#Ping 20.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

To test the communication with the correct MAC address:

R1(config)#int f0/1
R1(config-if)#mac-address 0000.1111.1111

R1#Ping 20.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

Note when IP source guard is enabled with source IP and MAC address Filtering, IP traffic is filtered based on the source IP and MAC address binding. The Switch forwards IP traffic when the source IP address of that traffic matches an entry in the DHCP snooping binding database or a manually created source binding table.

Task 19

Configure R4's F0/1 interface in VLAN 22 using the following parameters:

R4's F0/1 – 20.1.1.4 /24
Mac-address – 0000.4444.4444

On R4

R4(config)#int f0/1
R4(config-if)#ip address 20.1.1.4 255.255.255.0
R4(config-if)#mac-address 0000.4444.4444
R4(config-if)#no shut

On Cat-2

Cat-2(config)#int f0/4
Cat-2(config-if)#switchport mode access
Cat-2(config-if)#switchport access vlan 22

Cat-2(config-if)#spanning portfast

To test the configuration:

On R4

R4#Ping 20.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R4#Ping 20.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R4#Ping 20.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

Note R4 was added to VLAN 22 and was able to communicate with all the hosts/routers in VLAN 22.

Task 20

Configure DAI (Dynamic ARP Inspection) to fix the problem identified in the previous step such that if a new host/router is added to VLAN 22, it won't be able to communicate with any host/router in VLAN 22 unless it's IP to MAC address binding is added to the table.

On CAT-2

CAT-2(config)#ip arp inspection vlan 22
CAT-2(config)#ip arp inspection filter TST vlan 22 static

```
CAT-2(config)#arp access-list TST
CAT-2(config-arp-nacl)#permit ip host 20.1.1.1 mac host 0000.1111.1111
CAT-2(config-arp-nacl)#permit ip host 20.1.1.2 mac host 0000.2222.2222
CAT-2(config-arp-nacl)#permit ip host 20.1.1.3 mac host 0000.3333.3333
CAT-2(config-arp-nacl)#permit ip host 20.1.1.4 mac host 0000.4444.4444
```

To verify the configuration:

On R1

R1#Ping 20.1.1.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

R1#Ping 20.1.1.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

R1#Ping 20.1.1.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

To test the configuration:

On R4

```
R4(config)#int f0/1
R4(config-if)#no mac
```

R4#Ping 20.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note the MAC address does NOT match the binding in the Arp access-list.

R4(config)#int F0/1
R4(config-if)#mac-address 0000.4444.4444

R4#Ping 20.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

Note the MAC address is added to the F0/1 interface of R4 and the ping was successful.

To test by adding another router to this VLAN:

On R5

R5(config)#int F0/1
R5(config-if)#IP address 20.1.1.5 255.255.255.0
R5(config-if)#Mac-address 0000.5555.5555
R5(config-if)#no shut

R5#Ping 20.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 20.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 20.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 20.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

To allow for newly added routers/hosts to communicate with other hosts/routers in VLAN 22:

On Cat-2

Cat-2(config)#arp access-list TST
Cat-2(config-arp-nacl)#permit ip host 20.1.1.5 mac host 0000.5555.5555

To see the ARP access-list:

On Cat-2

Cat-2#Show arp access-list TST

ARP access list TST
    permit ip host 20.1.1.1 mac host 0000.1111.1111
    permit ip host 20.1.1.2 mac host 0000.2222.2222
    permit ip host 20.1.1.3 mac host 0000.3333.3333
    permit ip host 20.1.1.4 mac host 0000.4444.4444
    permit ip host 20.1.1.5 mac host 0000.5555.5555

To test the configuration:

On R5

R5#Ping 20.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R5#Ping 20.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:

```
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R5#Ping 20.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R5#Ping 20.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Task 21

Since the CPU of the switch performs the actual DAI validation checks, the incoming ARP packets should be configured to be 1/3 of its default value. This should ONLY be configured for the ports in VLAN 22.

Because the switch uses its CPU to perform Dynamic ARP Inspection, the switch will rate limit the number of ARP packets to 15 pps, this can be revealed using the "Show ip arp inspection interfaces" command, as follows:

On Cat-2

```
Cat-2#Sh ip arp inspection interfaces

 Interface        Trust State     Rate (pps)    Burst Interval
 ---------------  -----------     -----------   --------------
 Fa0/1            Untrusted            15              1
 Fa0/2            Untrusted            15              1
 Fa0/3            Untrusted            15              1
 Fa0/4            Untrusted            15              1

 Fa0/5            Untrusted            15              1
```
(The rest of the output is omitted)

Note the default value for all interfaces is set to 15 pps.

To configure the Cat to rate limit the number of ARP packets:

On Cat-2

```
Cat-2(config)#int range F0/1-5
Cat-2(config-if-range)#ip arp inspection limit rate 5 burst interval 1
```

To verify the configuration:

On Cat-2

```
Cat-2#Sh ip arp inspection interfaces

 Interface        Trust State     Rate (pps)    Burst Interval
 ---------------  -----------     ----------    --------------
 Fa0/1           Untrusted            5               1
 Fa0/2           Untrusted            5               1
 Fa0/3           Untrusted            5               1
 Fa0/4           Untrusted            5               1
 Fa0/5           Untrusted            5               1
```

(The rest of the output is omitted)

Task 22

Configure Cat-2 to keep track of all drop packets due to mismatch of the dynamic ARP inspection binding configured in one of the previous steps. The switch should log messages after 5 seconds of an event; ensure that the switch adds entries to the log buffer without generating a system message.

When the switch drops a packet, it adds an entry in the log buffer and generates a system message. Once the switch generates a system message, the particular entry is cleared from the log Buffer. The entry includes: VLAN, port number, Source and Destination IP and MAC addresses.

On Cat-2

```
Cat-2(config)#ip arp inspection log-buffer logs 0 interval 5
```

Note when the logs is set to 0, the switch will NOT generate a system message.

To test the configuration:

On R5

To test this configuration, the MAC address of R5 is removed and a ping is issued to emulate an invalid binding.

R5(config)#int f0/1
R5(config-if)#no mac

R5#Ping 20.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Cat-2#Sh ip arp inspection log

Total Log Buffer Size : 32
Syslog rate : 0 entries per 5 seconds.

```
Interface  Vlan  Sender MAC      Sender IP   Num Pkts   Reason    Time
---------- ----  --------------- ----------- ---------- --------- --------
Fa0/5       22   0012.d9d7.99a9  20.1.1.5            8  Acl Deny  21:18:15 UTC Tue Mar 2 1993
```

Task 23

Configure SNMP on Cat-1 using the following parameters:

- NMS IP address is 192.168.1.100
- RO community should be TST-RO
- RW community should be TST-RW
- The NMS is using Version 2C
- The community string should be "cisco"

On Cat-1

Cat-1(config)#snmp-server host 192.168.1.100 version 2c cisco
Cat-1(config)#snmp-server community TST-RO ro
Cat-1(config)#snmp-server community TST-RW rw

Task 24

Configure Cat-2 such that if in the future it is configured with BGP, it should send BGP notifications to the SNMP server with an address of 192.168.1.1. The switch should send these notifications using traps. Ensure that the switch uses version 2C and a community string of "cisco"

On Cat-2

Cat-2(config)#Snmp-server host 192.168.1.1 traps version 2C cisco bgp
Cat-2(config)#Snmp-server enable traps bgp

Task 25

Configure Cat-3 to send all traps to the host "PC1.MicronicsTraining.com" using community string of "cisco". The switch should resolve this FQDN to 10.1.1.200 locally.

On Cat-3

Cat-3(config)#**ip host PC1.MicronicsTraining.com** 10.1.1.200

Cat-3(config)#Snmp-server enable traps
Cat-3(config)#Snmp-server host PC1.MicronicsTraining.com cisco

Task 26

Configure SNMP on Cat-4 using the following parameters:
The SNMP manager must have Read-Only permission access to all objects using "cisco" as the string.

The switch should send VTP traps to 10.1.1.10, 10.1.1.100 and 10.1.1.200 using the following SNMP versions:
Host 10.1.1.10 SNMP version 1, Host 10.1.1.100 and 10.1.1.200 SNMPv2C
Ensure that the community string of "cisco" is sent with the traps:

On Cat-4

```
Cat-4(config)#Snmp-server community cisco
Cat-4(config)#Snmp-server enable traps VTP
Cat-4(config)#Snmp-server host 10.1.1.10 version 1 cisco
Cat-4(config)#Snmp-server host 10.1.1.100 version 2C cisco
Cat-4(config)#Snmp-server host 10.1.1.200 version 2C cisco
```

Task 27

Ensure that Cat-4 is configured with the following parameters for its previous SNMP configuration:

- Contact: Micronics Networking and Training Inc
- Location: Building A, Sydney office

On Cat-4

```
Cat-4(config)#snmp-s location Building A, Sydney office
Cat-4(config)#snmp-s contact Micronics Networking and Training Inc
```

Task 28

Configure Cat-1 such that whenever the switch learns or removes a MAC address on its port F0/18, an SNMP notification is generated and sent to the NMS located at 192.168.1.100. Since there are many users coming and going from the network, set up a trap interval time to bundle the notification traps and reduce network traffic using the following parameters:

- The traps should be generated every 30 minutes.
- The trap should contain a maximum of 150 entries.

This feature enables us to track users on a network by storing the Mac address activity on the switch. Once configured, every time a MAC address is learned or removed an SNMP notification is generated and sent to the NMS. On a very busy network when lots of users come and go, the default behavior is that an SNMP trap is sent every second. Because this can consume bandwidth, there are two parameters that can be configured to remedy this situation and they are as follows:

- **Mac address-table notification interval** – This value specifies the notification trap interval in seconds between each set of traps that are generated to the NMS. Default value is one second, and the range is 0 – 2,147,483,647 seconds.

- **Mac address-table notification history-size** – Specifies the maximum number of entries in the MAC notification history table. The default value is 1, and the range is 1 – 500 entries.

On Cat-1

Cat-1(config)#**Snmp-server host 192.168.1.100 traps private**

The above command identifies the NMS

Cat-1(config)#**Snmp-server enable traps mac-notification**

This command enables SNMP traps mac-notification

Cat-1(config)#**Mac-address-table notification**

The above command enables the mac address-table notification on the switch

Cat-1(config)#**Mac-address-table notification interval 1800**

This command sets the interval

Cat-1(config)#**Mac-address-table notification history-size 150**

This command sets the history-size

Cat-1(config)#Int f0/18
Cat-1(config-if)#**Snmp trap mac-notification** added

To enable the MAC notification trap whenever a MAC address is added

Cat-1(config-if)#**Snmp trap mac-notification removed**

To enable the MAC notification trap whenever a MAC address is removed

To verify the configuration:

On Cat-1

Cat-1#Show mac-address-table notification interface f0/18

MAC Notification Feature is Enabled on the switch

```
Interface          MAC Added Trap          MAC Removed Trap
---------          ------------------      ---------------------
FastEthernet0/18        Enabled                  Enabled
```

Cat-1#Show mac-address-table notification

MAC Notification Feature is Enabled on the switch
Interval between Notification Traps : 1800 secs
Number of MAC Addresses Added : 0
Number of MAC Addresses Removed : 0
Number of Notifications sent to NMS : 0
Maximum Number of entries configured in History Table : 150
Current History Table Length : 0
MAC Notification Traps are Enabled
History Table contents

Task 29

You received another request from your IT department to keep track of all the MAC addresses that are learned by Cat-2 port F0/18. The switch must use the NMS located at 192.168.1.1 /24, configure the switch to handle this request. You should use an IP address of 2.2.2.2 /8 to accomplish this task.

On Cat-2

Cat-2(config)#**Snmp-server host 192.168.1.1 trap private**
%IP_SNMP-3-SOCKET: can't open UDP socket
Unable to open socket on port 161

Note since this switch is not configured with an IP address, it will fail to configure the Snmp server. Therefore, an IP address should be configured before entering the "snmp-server" command as follows:

Cat-2(config)#Int lo0
Cat-2(config-if)#Ip addr 2.2.2.2 255.0.0.0

To setup the Snmp-Server:

Cat-2(config)#**snmp-server host 192.168.1.1 trap private**

Configures the switch to send mac-address traps to the NMS:

Cat-2(config)#**snmp-server enable traps mac-notification**

To enable MAC-address notification:

Cat-2(config)#**mac-address-table notification**

Cat-2(config)#Inter f0/18
Cat-2(config-if)#**snmp trap mac-notification** added

The above command enables the SNMP trap on interface F0/18 and configures the switch to send MAC notification traps whenever a MAC-address is added. If the switch must be configured to report the MAC addresses that are learnt and expired, then "snmp trap mac-notification removed" command must also be configured.

To verify the configuration:

On Cat-2

Cat-2#Show mac-address-table notification interface f0/18

MAC Notification Feature is Enabled on the switch

Interface	MAC Added Trap	MAC Removed Trap
FastEthernet0/18	Enabled	Disabled

Note if the "snmp trap mac-notification removed" command was also entered for F0/18 interface, under the "MAC removed Trap" column you will also see as "Enabled".

Cat-2#Show mac-address-table notification

MAC Notification Feature is Enabled on the switch
Interval between Notification Traps : 1 secs
Number of MAC Addresses Added : 0
Number of MAC Addresses Removed : 0
Number of Notifications sent to NMS : 0
Maximum Number of entries configured in History Table : 1
Current History Table Length : 0
MAC Notification Traps are Enabled
History Table contents

Task 30

Shut down the following ports:
The ports that connects Cat-3 to Cat-4
On Cat-1 and Cat-2 F0/23-24 and F0/21-22

On Cat-1

Cat-1(config)#int range F0/21-24
Cat-1(config-if-range)#shut

On Cat-2

Cat-2(config)#interface range F0/21-24
Cat-2(config-if-range)#Shut

On Cat-3

Cat-3(config)#Inter range F0/21-24
Cat-3(config-if-range)#Shut

On Cat-4

Cat-4(config)#Int range F0/21-24

```
Cat-4(config-if-range)#Shut
```

Task 31

Establish a trunk using an industry solution between Cat-1 and Cat-2 using ports F0/19-20; to STP these two ports should appear as one.
The ports on Cat-1 should be in passive negotiation state in which it should ONLY respond to PAgP packets.
Cat-2 should be configured appropriately.

PAgP is a Cisco proprietary protocol that can be used to automatically create EtherChannels by exchanging PAgP packets between Ethernet ports. PAgP has two modes of operation:

Auto: This mode places the port/s into a passive negotiation state, in which the ports ONLY respond to PAgP packets that they receive. Ports in this mode WILL NOT start PAgP packet negotiation, which minimizes the transmission of PAgP packets. If both ends of a given link are configured in AUTO mode, they will NOT negotiate a trunk.

Desirable: this mode places the port/s into an active negotiation state, in which the ports start negotiation by sending PAgP packets. Desirable mode will negotiate a trunk with another port configured in either AUTO or DESIRABLE mode.

On Cat-1

```
Cat-1(config)#int range f0/19-20
Cat-1(config-if-range)#swi trunk encap dot1q
Cat-1(config-if-range)#swi mode trunk
Cat-1(config-if-range)#channel-group 1 mode auto
```

On Cat-2

```
Cat-2(config)#int range f0/19-20
Cat-2(config-if-range)#swi trunk encap dot1q
Cat-2(config-if-range)#swi mode trunk
Cat-2(config-if-range)#channel-group 1 mode desirable
```

To verify the configuration:

On Cat-1

```
Cat-1#Sh int trunk

Port      Mode      Encapsulation  Status      Native vlan
Po1       on        802.1q         trunking    1
```
(The output is modified to only shows the Port Channel)

```
Cat-1#Show pagp neighbor

Flags:  S - Device is sending Slow hello.  C - Device is in Consistent state.
        A - Device is in Auto mode.        P - Device learns on physical port.

Channel group 1 neighbors
        Partner          Partner         Partner        Partner Group
Port    Name             Device ID       Port      Age  Flags   Cap.
Fa0/19  Cat-2            0019.2f90.ae00  Fa0/19    17s  SC      10001
Fa0/20  Cat-2            0019.2f90.ae00  Fa0/20    10s  SC      10001

Cat-1#Show etherchannel 1 summary

Flags:  D - down        P - in port-channel

        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+--------------------------
1      Po1(SU)        PAgP   Fa0/19(P)   Fa0/20(P)
```

Task 32

Configure the EtherChannel from the previous step such that packets sent to the same MAC address will use the same port.

EtherChannel load balancing can be configured in one of the following methods:

Source MAC address forwarding:
In this method, when the packets are sent to the EtherChannel they are distributed across the ports in the channel based on the source MAC address of the incoming packets. As a result of that, packets from different hosts use different ports.

Destination MAC address forwarding:
In this method, when the packets are sent to the EtherChannel they are distributed across the ports in the channel based on the destination host's MAC address, as a result of that, packets to the same destination, are forwarded out of the same port.

Source and Destination MAC address forwarding:
In this method, when the packets are sent to the EtherChannel they are distributed across the ports in the channel based on the both source and destination MAC address, as a result of that, packets from a given host to a given destination will use the same port.

Source IP address based forwarding:
In this method, when the packets are sent to the EtherChannel they are distributed across the ports in the channel based on the source IP address of the incoming packet. As a result of that, packets with different source IP address will use different port.

Destination IP address based forwarding:
In this method, when the packets are sent to the EtherChannel they are distributed across the ports in the channel based on the destination IP address of the incoming packet, as a result of that, packets to the same destination will use the same port.

Source and Destination IP address based forwarding:
In this method, when the packets are sent to the EtherChannel they are distributed across the ports in the channel based on the source and destination IP addresses. As a result of that, packets from a given IP source to a specific IP destination will use the same port.

On Both Switches:

(config)#Port-channel load-balance dst-mac

To verify the configuration:

On Cat-1

Cat-1#Show etherchannel load-balance

EtherChannel Load-Balancing Operational State (dst-mac):
Non-IP: Destination MAC address

IPv4: Destination MAC address
IPv6: Destination IP address

Task 33

Configure a layer 3 EtherChannel using ports F0/21-22 connecting Cat-2 to Cat-4. These ports should NOT use any protocol/s to negotiate an EtherChannel. Use the following IP addresses:
Cat-2 – 10.1.24.2 /24 and Cat-4 – 10.1.24.4 /24

When configuring a layer 3 EtherChannels the port-channel interface should be created first and then, assigned to the physical port using the "Channel-group" command.

On Cat-2

```
Cat-2(config)#int port-channel 24
Cat-2(config-if)#NO switchport
Cat-2(config-if)#ip address 10.1.24.2 255.255.255.0

Cat-2(config)#int range F0/21-22
Cat-2(config-if-range)#no switchport
Cat-2(config-if-range)#channel-group 24 mode on
Cat-2(config-if-range)#NO shut
```

On Cat-4

Before configuring this switch for a layer 3 EtherChannel, remember that you must change the "SDM prefer VLAN", or else the IOS will NOT allow you to create a port-channel interface.

```
Cat-4(config)sdm prefer routing
Cat-4#reload
```

After the switch is reloaded:

```
Cat-4(config)#Int port-channel 24
Cat-4(config-if)#NO swi
Cat-4(config-if)#ip address 10.1.24.4 255.255.255.0
```

```
Cat-4(config)#int range f0/21-22
Cat-4(config-if-range)#NO swi
Cat-4(config-if-range)#channel-group 24 mode on
Cat-4(config-if-range)#no shut
```

Note in this case we must use the "Channel-group 24 mode ON" command, the ON tells the switch NOT to use PAgP or LACP to negotiate the EtherChannel. With mode both ends of the links should be configured with mode set to "ON".

To test the configuration:

On Cat-2

Cat-2#Ping 10.1.24.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.24.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

To verify the configuration:

Cat-2

Cat-2#Show etherchannel summary

```
Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 2
Number of aggregators:           2

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------
1      Po1(SU)         PAgP      Fa0/19(P)   Fa0/20(P)
24     Po24(RU)         -        Fa0/21(P)   Fa0/22(P)
```

Note the letter "R" to the right of the Po24 states that this is a layer 3 EtherChannel,

whereas, the letter "S" to the right of the Po1, states that the link is a layer 2 EtherChannel.

Task 34

Establish a trunk using Cisco proprietary solution between Cat-4 and Cat-3 using ports F0/19-20; to STP these two ports should appear as one.
The ports on Cat-4 should be configured such that they start negotiation process by sending LACP packets. The ports on Cat-3 should NOT be configured the same.

LACP is an industry standard (IEEE 802.3ad) solution for managing EtherChannels between the switches. LACP offers two modes of operation:

Active: In this mode the ports are placed into an active negotiation state, in which the ports involved start negotiating with other ports by sending LACP packets. If both ends of a given link are configured in Active or Passive mode, the ports will negotiate an EtherChannel.

Passive: In this mode the ports are placed into a passive mode, in which the ports can ONLY respond to LACP packets that they receive. If both ends of a given link are configured in Passive mode, the ports will NOT negotiate an EtherChannel, whereas, an active mode configured on one side and Passive configured on the other the switches will negotiate an EtherChannel link.

On Cat-4

```
Cat-4(config)#int range f0/19-20
Cat-4(config-if-range)#swi trunk encap isl
Cat-4(config-if-range)#swi mode trunk
Cat-4(config-if-range)#channel-group 34 mode active
Cat-4(config-if-range)#no shut
```

On Cat-3

```
Cat-3(config)#int range F0/19-20
Cat-3(config-if-range)#swi trunk encap isl
Cat-3(config-if-range)#swi mode trunk
Cat-3(config-if-range)#channel-group 34 mode passive
Cat-3(config-if-range)#NO shut
```

To verify the configuration:

On Cat-3

Cat-3#Sh int trunk

```
Port      Mode        Encapsulation  Status      Native vlan
Po34       on          isl            trunking     1

Port      Vlans allowed on trunk
Po34       1-4094

Port      Vlans allowed and active in management domain
Po34       1

Port      Vlans in spanning tree forwarding state and not pruned
Po34       1
```

Cat-3#Sh etherchannel summ

```
Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
34     Po34(SU)        LACP      Fa0/19(P)   Fa0/20(P)
```

Task 35

In the future you will be adding another 14 ports to this EtherChannel, ensure that port F0/19 of Cat-3 and Cat-4 will be one of the ports that will be in active state first and not standby.

Before any changes are made, the default parameters should be checked, as follows:

Cat-3#Sh lacp 34 internal

```
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode

Channel group 34
                         LACP port    Admin   Oper   Port     Port
Port      Flags  State   Priority     Key     Key    Number   State
Fa0/19    SP     bndl    32768        0x22    0x22   0xF      0x3C
Fa0/20    SP     bndl    32768        0x22    0x22   0x10     0x3C
```

When LACP is configured, it will try to use maximum number of ports in a given channel, up to a maximum of 16 ports. But only 8 ports can be active at any time, the additional ports are placed in a hot-standby state, this decision is made by the system; if one of the active ports goes down, one of the hot-standby links will become active.
Every link has a unique priority which is made up of:

- **LACP system priority**
- **System-ID (Which is a combination of LACP-Priority and switch MAC address)**
- **LACP port priority**
- **Port number**

Numerically lower value will always have a higher priority.
This priority decides which ports should be place in hot-standby mode, and which ports should be in Active mode.

On Both Switches

(config)#int f0/19
(config-if)#**lacp port-priority 1**

To verify the configuration:

On Cat-3

Cat-3#Sh lacp 34 internal

```
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode
Channel group 34
```

Port	Flags	State	LACP port Priority	Admin Key	Oper Key	Port Number	Port State
Fa0/19	SP	bndl	1	0x22	0x22	0xF	0x3C
Fa0/20	SP	bndl	32768	0x22	0x22	0x10	0x3C

Task 36

Configure a layer 3 EtherChannel using ports F0/21-22 connecting Cat-1 to Cat-3. These ports should use IEEE 802.ad to negotiate an EtherChannel. Use the following IP addresses:
Cat-1 – 10.1.13.1 /24 and Cat-3 – 10.1.13.3 /24

On Cat-1

```
Cat-1(config)#int port-channel 13
Cat-1(config-if)#no swi
Cat-1(config-if)#ip addr 10.1.13.1 255.255.255.0

Cat-1(config-if)#int range f0/21-22
Cat-1(config-if-range)#no swi
Cat-1(config-if-range)#channel-group 13 mode passive
Cat-1(config-if-range)#No shut
```

On Cat-3

```
Cat-3(config)#int port-channel 13
Cat-3(config-if)#no swi
Cat-3(config-if)#ip address 10.1.13.3 255.255.255.0

Cat-3(config-if)#int range f0/21-22
Cat-3(config-if-range)#no swi
Cat-3(config-if-range)#channel-group 13 mode active
Cat-3(config-if-range)#no Shut
```

To verify the configuration:

On Cat-1

```
Cat-1#Ping 10.1.13.3
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Task 37

Erase the config.text and vlan.dat file and reload the switches before proceeding to the next task.

On All Switches:

```
Cat-1#pwd
This command display the current working directory

flash:

#delete config.text
Delete filename [config.text]?
Delete flash:config.text? [confirm]

#delete vlan.dat
Delete filename [vlan.dat]?
Delete flash:vlan.dat? [confirm]

Cat-1#dir
Directory of flash:/

    4  -rwx     7252875   Mar 1 1993 00:03:37 +00:00  c3560-advipservicesk9-mz.122-25.SEB4.bin
    5  drwx         192   Mar 1 1993 00:05:36 +00:00  c3560-ipbase-mz.122-25.SEB4
              15998976 bytes total (1731072 bytes free)
```

Task 38

Configure a trunk between Cat-1 and Cat-2 using ports F0/19-20; use a Cisco proprietary trunking solution to accomplish this task. You should Shutdown ports F0/21-24 on both Cat-2 and Cat-2.

On Both Switches:

```
(config)#int range f0/19-20
(config-if-range)#swi trunk encap isl
(config-if-range)#swi mode trunk

(config)#int range F0/21-24
(config-if-range)#shut
```

Task 39

Configure R1 and R2 based on the following parameters and ensure that the following ports are in VLAN 12:

R1's F0/0 – 10.1.12.1 /24 and R2's F0/1 – 10.1.12.2 /24

On R1

```
R1(config)#int f0/0
R1(config-if)#ip addr 10.1.12.1 255.255.255.0
R1(config-if)#no shut
```

On R2

```
R2(config)#int f0/1
R2(config-if)#ip addr 10.1.12.2 255.255.255.0
R2(config-if)#NO shut
```

On Cat-1

```
Cat-1(config)#vtp domain TST

Cat-1(config)#int f0/1
Cat-1(config-if)#swi mode acc

Cat-1(config-if)#swi acc v 12
Cat-1(config-if)#spanning portfast
```

On Cat-2

```
Cat-2(config)#int f0/2
Cat-2(config-if)#swi mode acc
Cat-2(config-if)#swi acc v 12
Cat-2(config-if)#spanning portfast
```

To test the configuration:

On R1

R1#Ping 10.1.12.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Task 40

Configure Cat-1 such that it marks all traffic from R1 with an IP Precedence of 1.

On Cat-1

```
Cat-1(config)#mls qos
```

QOS should be enabled, if QOS is NOT enabled, the policy will NOT have any affect.

```
Cat-1(config)#Access-list 100 permit ip any any

Cat-1(config)#class-map QOS
Cat-1(config-cmap)#match access-group 100

Cat-1(config-cmap)#policy-map TST
Cat-1(config-pmap)#class QOS
Cat-1(config-pmap-c)#set ip precedence 1

Cat-1(config-pmap-c)#int f0/1
Cat-1(config-if)#service-policy input TST
```

Note on 3560s there are few things that are NOT supported and they are:

Service-policy is NOT supported on the Outbound direction, you should receive the following message:

Warning: Assigning a policy map to the output side of an interface not supported

In the class-map the "input-interface" can NOT be used, if it is used you will get the following message when applying the "Service-policy" to an interface:

%QoS: policy-map TST with MATCH INPUT-INTERFACE not allowed on non-SVI interface Service Policy attachment failed

To verify the configuration:

On Cat-1

Cat-1#Sh class-map

```
Class Map match-any class-default (id 0)
  Match any

 Class Map match-all QOS (id 1)
  Match access-group  100
```

Cat-1#Show access-list

```
Extended IP access list 100
   10 permit ip any any
```

Cat-1#Show policy-map

```
  Policy Map TST
    Class QOS
     set ip precedence 1
```

To test the configuration:

To test the configuration, an access-list should be created permitting each IP Precedence value with a log option. This is created so we can test different traffic marked with different IP Precedence levels generated by R1.

On R2

R2(config)#access-list 100 permit ip any any Precedence 0 log

```
R2(config)#access-list 100 permit ip any any Precedence 1 log
R2(config)#access-list 100 permit ip any any Precedence 2 log
R2(config)#access-list 100 permit ip any any Precedence 3 log
R2(config)#access-list 100 permit ip any any Precedence 4 log
R2(config)#access-list 100 permit ip any any Precedence 5 log
R2(config)#access-list 100 permit ip any any Precedence 6 log
R2(config)#access-list 100 permit ip any any Precedence 7 log
R2(config)#access-list 100 permit ip any any log

R2(config)#int f0/1
R2(config-if)#ip access-group 100 in
```

To test the configuration:

Generate traffic from R1:

On R1

R1#Ping 10.1.12.2 repeat 10

```
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 1/2/4 ms
```

On R2

R2#Sh ip access-list 100

```
Extended IP access list 100
    10 permit ip any any precedence routine log
    20 permit ip any any precedence priority log (10 matches)
    30 permit ip any any precedence immediate log
    40 permit ip any any precedence flash log
    50 permit ip any any precedence flash-override log
    60 permit ip any any precedence critical log
    70 permit ip any any precedence internet log
    80 permit ip any any precedence network log
    90 permit ip any any log
```

Note the 10 ICMP packets matched IP Precedence 1. The reason that the packets inbound to R1 have preserved their marking is because the QOS on the second switch (Cat-2) is disabled. If the "MLS QOS" is disabled, the packets will traverse through the switch with their marking untouched. If the "MLS QOS" is enabled, the switch will remark all packets

with IP Precedence of 0. To test this, the QOS of the second switch should be enabled as follows:

On Cat-2

Cat-1(config)#Mls qos

To verify the configuration:

On Cat-2

Cat-2#Show mls qos

QoS is enabled
QoS ip packet dscp rewrite is enabled

To generate some traffic on R1:

Note 10 pings are initiated

R1#Ping 10.1.12.2 repeat 10

```
Type escape sequence to abort.
Sending 15, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 1/1/4 ms
```

To verify the configuration:

On R2

R2#Sh access-list

```
Extended IP access list 100
    10 permit ip any any precedence routine log (10 matches)
    20 permit ip any any precedence priority log (10 matches)
    30 permit ip any any precedence immediate log
    40 permit ip any any precedence flash log
    50 permit ip any any precedence flash-override log
    60 permit ip any any precedence critical log
    70 permit ip any any precedence internet log
    80 permit ip any any precedence network log
    90 permit ip any any log
```

Note when traffic from R1 traversed through Cat-1, Cat-1 remarked the traffic with IP

Precedence of 1, but because QOS was enabled on another switch along the path to R2, in this case Cat-2, when the traffic traversed that switch, the switch remarked the traffic back to zero.

Task 41

Ensure that the traffic from R1 retains its Precedence level; DO NOT disable QOS on Cat-2.

On Cat-2

Cat-2(config)#int range f0/19-20
Cat-2(config-if-range)#**mls qos trust ip-precedence**

The above command shows how to set the trusted state of an interface to IP precedence

To test the configuration:

On R2

R2#Clear access-list counters

On R1

R1#Ping 10.1.12.2 repeat 25

Type escape sequence to abort.
Sending 25, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (25/25), round-trip min/avg/max = 1/1/4 ms

On R2

R2#Sh access-list

Extended IP access list 100
    10 permit ip any any precedence routine log
    20 permit ip any any precedence priority log (25 matches)
    30 permit ip any any precedence immediate log

```
40 permit ip any any precedence flash log
50 permit ip any any precedence flash-override log
60 permit ip any any precedence critical log
70 permit ip any any precedence internet log
80 permit ip any any precedence network log
90 permit ip any any log
```

Note the IP Precedence level is retained.

Task 42

Configure R3 and R4 in VLAN 34 and R5 and R6 in VLAN 56; use the following IP addresses to configure these routers.

R3's F0/0 – 10.1.34.3 /24 and R4's F0/1 – 10.1.34.4 /24
R5's F0/0 – 10.1.56.5 /24 and R6's F0/1 – 10.1.56.6 /24

You should provide Inter-VLAN routing between these two VLANs, and Vlan 12, use the following IP addresses as their default gateway:
For Vlan 12 – 10.1.12.100 /24
For Vlan 34 – 10.1.34.100 /24
For Vlan 56 – 10.1.56.100 /24

On R3

```
R3(config)#int f0/0
R3(config-if)#ip addr 10.1.34.3 255.255.255.0
R3(config-if)#no shut
```

On R4

```
R4(config)#int f0/1
R4(config-if)#ip address 10.1.34.4 255.255.255.0
R4(config-if)#no shut
```

On R5

```
R5(config)#int f0/0
R5(config-if)#ip addr 10.1.56.5 255.255.255.0
R5(config-if)#no shut
```

On R6

```
R6(config)#int f0/1
R6(config-if)#ip address 10.1.56.6 255.255.255.0
R6(config-if)#no shut
```

On Cat-1

```
Cat-1(config)#int f0/3
Cat-1(config-if)#Switchport mode access
Cat-1(config-if)#switchport access vlan 34
Cat-1(config-if)#spanning portfast

Cat-1(config)#int f0/5
Cat-1(config-if)#Switchport mode access
Cat-1(config-if)#switchport access vlan 56
Cat-1(config-if)#spanning portfast
```

On Cat-2

```
Cat-2(config)#int f0/4
Cat-2(config-if)#swi mode access
Cat-2(config-if)#swi access vlan 34
Cat-2(config-if)#spanning portfast

Cat-2(config)#int f0/6
Cat-2(config-if)#swi mode access
Cat-2(config-if)#swi access vlan 56
Cat-2(config-if)#spanning portfast
```

To provide Inter-Vlan routing:

On Cat-2

```
Cat-2(config)#Ip routing
Cat-2(config)#int vlan 12
Cat-2(config-if)#ip address 10.1.12.100 255.255.255.0

Cat-2(config)#int vlan 34
Cat-2(config-if)#ip addr 10.1.34.100 255.255.255.0

Cat-2(config)#int vlan 56
Cat-2(config-if)#ip addr 10.1.56.100 255.255.255.0
```

On R1 and R2

(config)# ip route 0.0.0.0 0.0.0.0 10.1.12.100

On R3 and R4

(config)# ip route 0.0.0.0 0.0.0.0 10.1.34.100

On R5 and R6

(config)# ip route 0.0.0.0 0.0.0.0 10.1.56.100

To test the configuration:

On R1

R1#Ping 10.1.34.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 10.1.34.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 10.1.56.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.56.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 10.1.56.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.56.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Task 43

Ensure that the traffic in VLAN 12 is marked with Precedence 3 and the traffic in VLAN 34 is marked with Precedence 4. All other traffic should be set to Precedence 0.

To configure this task, two Class-maps are created, one is called VLAN-12 and the second one is called VLAN-34, then, two policy-maps are created, one called VLAN-12 and the second one is called VLAN-34, then, the policy-map VLAN-12 is applied to interface Vlan 12 and policy-map VLAN-34 is applied to interface vlan 34, lastly, the "mls qos vlan-based" command is applied to the physical interfaces in the trunk.

On Cat-2

```
Cat-2(config)#access-list 100 permit ip any any

Cat-2(config)#class-map VLAN-12
Cat-2(config-cmap)#match access-group 100

Cat-2(config)#Policy-map VLAN-12
Cat-2(config-pmap)#Class VLAN-12
Cat-2(config-pmap-c)#set ip precedence 3

Cat-2(config)#Class-map VLAN-34
Cat-2(config-cmap)#match access-group 100

Cat-2(config)#Policy-map VLAN-34
Cat-2(config-pmap)#Class VLAN-34
Cat-2(config-pmap-c)#set ip precedence 4

Cat-2(config)#interface Vlan 12
Cat-2(config-if)#service-policy in VLAN-12

Cat-2(config)#interface vlan 34
Cat-2(config-if)#service-policy in VLAN-34

Cat-2(config)#int range f0/19-20

Cat-2(config-if-range)#mls qos vlan-based
```

To test the configuration:

On R2

```
R2#Clear access-list counters
```

This command is entered to clear the counters on configured access-list

On R1

R1#Ping 10.1.12.2 repeat 12

Type escape sequence to abort.
Sending 12, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!!!!!!!!
Success rate is 100 percent (12/12), round-trip min/avg/max = 1/2/4 ms

The above Ping is repeated 12 times so it could be identified as traffic coming from VLAN 12

On R3

R3#Ping 10.1.12.2 repeat 34

Type escape sequence to abort.
Sending 34, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (34/34), round-trip min/avg/max = 1/2/4 ms

The above Ping is repeated 34 times so it could be identified as traffic coming from VLAN 34

On R5

R5#Ping 10.1.12.2 repeat 56

Type escape sequence to abort.
Sending 56, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!
Success rate is 100 percent (56/56), round-trip min/avg/max = 1/1/4 ms

The above Ping is repeated 56 times so it could be identified as traffic coming from VLAN 56

To verify the configuration:

On R2

```
R2#Sh access-list

Extended IP access list 100
    10 permit ip any any precedence routine log (56 matches)
    20 permit ip any any precedence priority log
    30 permit ip any any precedence immediate log
    40 permit ip any any precedence flash log (12 matches)
    50 permit ip any any precedence flash-override log (34 matches)
    60 permit ip any any precedence critical log
    70 permit ip any any precedence internet log
    80 permit ip any any precedence network log
    90 permit ip any any log
```

Untagged traffic (VLAN 56)

Traffic from VLAN 34 tagged with IP Precedence 4

Traffic from VLAN 12 tagged with IP Precedence 3

Task 44

Erase the config.text and Vlan.dat and reload the switches before proceeding to the next lab.

Lab 5 – Advanced Spanning-tree protocol Configuration

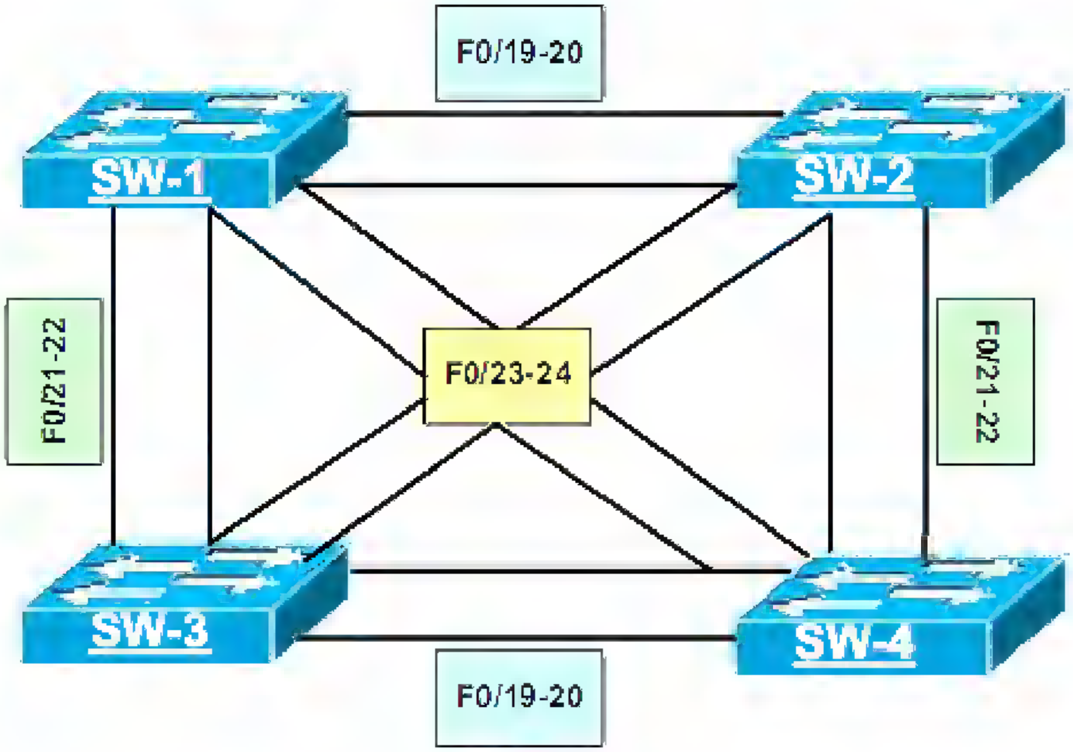

Task 1

Shut down all ports except ports F0/19-22 on all switches.

On All Switches:

```
(config)#Interface range F0/1-18 , F0/23-24
(config-if-range)#Shutdown
```

To verify the configuration:

On All Switch:

```
Port      Name        Status     Vlan      Duplex  Speed  Type
Fa0/1                 disabled   1           auto   auto  10/100BaseTX
Fa0/2                 disabled   1           auto   auto  10/100BaseTX
```

```
Fa0/3                    disabled    1              auto    auto   10/100BaseTX
Fa0/4                    disabled    1              auto    auto 10/100BaseTX
Fa0/5                    disabled    1              auto    auto 10/100BaseTX
Fa0/6                    disabled    1              auto    auto 10/100BaseTX
Fa0/7                    disabled    1              auto    auto 10/100BaseTX
Fa0/8                    disabled    1              auto    auto 10/100BaseTX
Fa0/9                    disabled    1              auto    auto 10/100BaseTX
Fa0/10                   disabled    1              auto    auto 10/100BaseTX
Fa0/11                   disabled    1              auto    auto 10/100BaseTX
Fa0/12                   disabled    1              auto    auto 10/100BaseTX
Fa0/13                   disabled    1              auto    auto 10/100BaseTX
Fa0/14                   disabled    1              auto    auto 10/100BaseTX
Fa0/15                   disabled    1              auto    auto 10/100BaseTX
Fa0/16                   disabled    1              auto    auto 10/100BaseTX
Fa0/17                   disabled    1              auto    auto 10/100BaseTX
Fa0/18                   disabled    1              auto    auto 10/100BaseTX
Fa0/19                  connected 1            a-full  a-100 10/100BaseTX
Fa0/20                  connected 1            a-full  a-100 10/100BaseTX
Fa0/21                  connected trunk        a-full  a-100 10/100BaseTX
Fa0/22                  connected trunk        a-full  a-100 10/100BaseTX
Fa0/23                   disabled   1              auto    auto 10/100BaseTX
Fa0/24                   disabled   1              auto    auto 10/100BaseTX
```

Task 2

Configure ports F0/19-20 between SW-1 and SW2, and between SW-3 and SW-4 as two trunk ports; you should use an industry standard protocol to accomplish this task. These ports should never become an access port through negotiation.

On SW-1 and SW-2

```
(config)#int range f0/19-20
(config-if-range)#Switch trunk encap dot1q
(config-if-range)#Switch mode trunk
```

To verify the configuration:

On SW-1

SW-1#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1

(The rest of the output is omitted)

On SW-2

SW-2#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1

(The rest of the output is omitted)

On SW-3 and SW-4

```
(config)#int range f0/19-20
(config-if-range)#Switch trunk encap dot1q
(config-if-range)#Switch mode trunk
```

To verify the configuration:

On SW-3

SW-3#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1

(The rest of the output is omitted)

On SW-4

SW-4#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1

(The rest of the output is omitted)

Task 3

Configure ports F0/21-22 between SW-2 and SW4, and between SW-1 and SW-3 as two trunk ports; you should use an industry standard protocol to accomplish this task. These ports should never become an access port through negotiation.

On SW-2 and SW-4

```
(config)#int range f0/21-22
(config-if-range)#Switch trunk encap dot1q
(config-if-range)#Switch mode trunk
```

To verify the configuration:

On SW-4

SW-4#Show int trunk

```
Port     Mode      Encapsulation  Status     Native vlan
Fa0/19    on        802.1q        trunking    1
Fa0/20    on        802.1q        trunking    1
Fa0/21    on        802.1q        trunking    1
Fa0/22    on        802.1q        trunking    1
```

(The rest of the output is omitted)

On SW-2

SW-2#Show int trunk

```
Port     Mode      Encapsulation  Status     Native vlan
Fa0/19    on        802.1q        trunking    1
Fa0/20    on        802.1q        trunking    1
Fa0/21    on        802.1q        trunking    1
Fa0/22    on        802.1q        trunking    1
```

(The rest of the output is omitted)

On SW-1 and SW-3

```
(config)#int range f0/21-22
(config-if-range)#Switch trunk encap dot1q
(config-if-range)#Switch mode trunk
```

To verify the configuration:

On SW-1

SW-1#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1
Fa0/21	on	802.1q	trunking	1
Fa0/22	on	802.1q	trunking	1

(The rest of the output is omitted)

On SW-3

SW-3#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1
Fa0/21	on	802.1q	trunking	1
Fa0/22	on	802.1q	trunking	1

(The rest of the output is omitted)

Task 4

These switches should be configured in a VTP domain called "CCIE".

On SW-1

(config)#vtp domain CCIE

This configuration will be propagated via VTP to the other switches.

To verify the configuration:

On SW-3

SW-3#Show vtp status | inc VTP Domain Name

VTP Domain Name : CCIE

Note the domain name is propagated by VTP

Task 5

Create the following VLANs and ensure that they are propagated to all four switches: 100, 200, 300, 400, 500 and 600

On SW-1

```
SW-1(config)#vlan 100,200,300,400,500,600
SW-1(config-vlan)#exit
```

To verify the configuration:

On SW-1

SW-1#Sh vlan br | exc unsup

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
100  VLAN0100                         active
200  VLAN0200                         active
300  VLAN0300                         active
400  VLAN0400                         active
500  VLAN0500                         active
600  VLAN0600                         active
```

On SW-4

SW-4#Sh vlan br | exc unsup

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
```

```
                                            Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                            Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                            Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                            Fa0/17, Fa0/18, Fa0/23, Fa0/24
                                            Gi0/1, Gi0/2
100  VLAN0100                         active
200  VLAN0200                         active
300  VLAN0300                         active
400  VLAN0400                         active
500  VLAN0500                         active
600  VLAN0600                         active
```

Task 6

Ensure that SW-1 is the root bridge for VLAN 100, SW-2 is the root bridge for VLAN 200, SW-3 is the root bridge for VLAN 300 and SW-4 is the root bridge for VLAN 400. You should use a macro to accomplish this task.

On SW-1

SW-1(config)#spanning-tree vlan 100 root primary

On SW-2

SW-2(config)#Spanning-tree vlan 200 root primary

On SW-3

SW-3(config)#Spanning-tree vlan 300 root primary

On SW-4

SW-4(config)#Spanning-tree vlan 400 root primary

To verify the configuration:

On SW-1

SW-1#Sh spanning-tree VLAN 100

```
VLAN0100
  Spanning tree enabled protocol ieee
  Root ID    Priority    24676
             Address     001b.2be5.0e00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24676  (priority 24576 sys-id-ext 100)

             Address     001b.2be5.0e00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 15

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Desg FWD 19        128.23   P2p
Fa0/22           Desg FWD 19        128.24   P2p
```

On SW-2

SW-2#Sh spanning-tree VLAN 200

```
VLAN0200
  Spanning tree enabled protocol ieee
  Root ID    Priority    24776
             Address     001c.575f.fd00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24776  (priority 24576 sys-id-ext 200)
             Address     001c.575f.fd00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 15

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Desg FWD 19        128.23   P2p
Fa0/22           Desg FWD 19        128.24   P2p
```

On SW-3

SW-3#Sh spanning-tree VLAN 300

```
VLAN0300
  Spanning tree enabled protocol ieee
  Root ID    Priority    24876
             Address     000d.65ea.3180  <------------ Note this matches the MAC of this
             This bridge is the root                   Switch
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24876  (priority 24576 sys-id-ext 300)
             Address     000d.65ea.3180  <------------
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 15

Interface        Role Sts   Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Desg FWD 19        128.23   P2p
Fa0/22           Desg FWD 19        128.24   P2p
```

On SW-4

SW-4#Sh spanning-tree VLAN 400

```
VLAN0400
  Spanning tree enabled protocol ieee
  Root ID    Priority    24976
             Address     000d.65e1.9200
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24976  (priority 24576 sys-id-ext 400)
             Address     000d.65e1.9200
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts   Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Desg FWD 19        128.23   P2p
Fa0/22           Desg FWD 19        128.24   P2p
```

Task 7

Implement the following policy:

1. VLAN 100 should never traverse SW-4
2. VLAN 200 should never traverse SW-3
3. VLAN 300 should never traverse SW-2
4. VLAN 400 should never traverse SW-1

The first Policy

On SW-2

```
SW-2(config)#int range f0/21-22
SW-2(config-if-range)#switchport trunk allowed vlan except 100
```

On SW-3

```
SW-3(config)#int range f0/19-20
SW-3(config-if-range)#switchport trunk allowed vlan except 100
```

On SW-4

```
SW-4(config)#int range f0/19-22
SW-4(config-if-range)#switchport trunk allowed vlan except 100
```

To verify the configuration:

On SW-2

SW-2#Sh int trunk

```
Port      Mode      Encapsulation  Status      Native vlan
Fa0/19    on        802.1q         trunking    1
Fa0/20    on        802.1q         trunking    1
Fa0/21    on        802.1q         trunking    1
Fa0/22    on        802.1q         trunking    1

Port      Vlans allowed on trunk
Fa0/19    1-4094
Fa0/20    1-4094
Fa0/21    1-99,101-4094  <-------  Note VLAN 100 is NOT allowed on the
Fa0/22    1-99,101-4094  <-------  trunk
```

```
Port        Vlans allowed and active in management domain
Fa0/19      1,100,200,300,400,500,600
Fa0/20      1,100,200,300,400,500,600
Fa0/21      1,200,300,400,500,600
Fa0/22      1,200,300,400,500,600

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,100,200,400,500,600
Fa0/20      1,200,400,500,600
Fa0/21      1,200,300,400,500,600

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      200
```

On SW-3

SW-3#Show int trunk

```
Port     Mode     Encapsulation  Status      Native vlan
Fa0/19   on       802.1q         trunking    1
Fa0/20   on       802.1q         trunking    1
Fa0/21   on       802.1q         trunking    1
Fa0/22   on       802.1q         trunking    1

Port        Vlans allowed on trunk
Fa0/19      1-99,101-4094
Fa0/20      1-99,101-4094
Fa0/21      1-4094
Fa0/22      1-4094
```

Note VLAN 100 is NOT allowed on the trunk

```
Port        Vlans allowed and active in management domain
Fa0/19      1,200,300,400,500,600
Fa0/20      1,200,300,400,500,600
Fa0/21      1,100,200,300,400,500,600
Fa0/22      1,100,200,300,400,500,600

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,200,300,400,500,600
Fa0/20      300
Fa0/21      1,100,300,400,500,600

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      1,300,400,500,600
```

On SW-4

SW-4#Sho int trunk

```
Port      Mode   Encapsulation  Status      Native vlan
Fa0/19    on     802.1q         trunking    1
Fa0/20    on     802.1q         trunking    1
Fa0/21    on     802.1q         trunking    1
Fa0/22    on     802.1q         trunking    1

Port      Vlans allowed on trunk
Fa0/19    1-99,101-4094
Fa0/20    1-99,101-4094
Fa0/21    1-99,101-4094
Fa0/22    1-99,101-4094

Port      Vlans allowed and active in management domain
Fa0/19    1,200,300,400,500,600
Fa0/20    1,200,300,400,500,600
Fa0/21    1,200,300,400,500,600
Fa0/22    1,200,300,400,500,600

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/19    1,300,400,500,600
Fa0/20    1,400,500,600
Fa0/21    1,200,300,400,500,600

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/22    1,300,400,500,600
```

Note VLAN 100 is NOT allowed on the trunk.

The second policy:

On SW-1

```
SW-1(config)#int range f0/21-22
SW-1(config-if-range)#Switch trunk allowed vlan except 200
```

To verify the configuration:

On SW-1

SW-1#Sh int trunk

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1
Fa0/21      on           802.1q         trunking      1
Fa0/22      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/20      1-4094
Fa0/21      1-199,201-4094
Fa0/22      1-199,201-4094
```

Note VLAN 200 is NOT allowed to traverse the trunk links connecting this switch to SW-3

On SW-4

```
SW-4(config)#int range f0/19-20
SW-4(config-if-range)#Switch trunk allowed vlan except 100,200
```

To verify the configuration:

SW-4

SW-4#Sh int trunk

```
Port        Mode       Encapsulation  Status        Native vlan
Fa0/19      on         802.1q         trunking      1
Fa0/20      on         802.1q         trunking      1
Fa0/21      on         802.1q         trunking      1
Fa0/22      on         802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-99,101-199,201-4094
Fa0/20      1-99,101-199,201-4094
Fa0/21      1-99,101-4094
Fa0/22      1-99,101-4094
```

Note VLAN 200 is NOT allowed on the trunk ports connecting this switch to SW-3

On SW3

The following configuration on ports F0/19-20 has to be reconfigure to deny

VLAN 100 again, or else, the command for VLAN 200 will override the previous configuration that was denying VLAN 100.

```
SW-3(config)#int range f0/19-20
SW-3(config-if-range)#switchport trunk allowed vlan except 100,200

SW-3(config)#int range f0/21-22
SW-3(config-if-range)#switchport trunk allowed vlan except 200
```

The third Policy:

On SW-2

```
SW-2(config-if-range)#int range f0/21-22
SW-2(config-if-range)#swi trunk allowed vlan except 100,300

SW-2(config-if-range)#int range f0/19-20
SW-2(config-if-range)#swi trunk allowed vlan except 300
```

To verify the configuration:

On SW-2

```
SW-2#Sh int trunk

Port      Mode        Encapsulation  Status      Native vlan
Fa0/19    on          802.1q         trunking    1
Fa0/20    on          802.1q         trunking    1
Fa0/21    on          802.1q         trunking    1
Fa0/22    on          802.1q         trunking    1

Port      Vlans allowed on trunk
Fa0/19     1-299,301-4094
Fa0/20     1-299,301-4094
Fa0/21     1-99,101-299,301-4094
Fa0/22     1-99,101-299,301-4094
```

Note the above output show that on ports F0/19-20 ONLY VLAN 300 is denied, whereas, on ports F0/21-22, VLANs 100 and 300 are denied.

On SW-4

```
SW-4(config)#int range f0/21-22
SW-4(config-if-range)#swi trunk allowed vlan except 100,300
```

To verify the configuration:

On SW-4

SW-4#Sh int trunk

```
Port      Mode   Encapsulation  Status     Native vlan
Fa0/19    on     802.1q         trunking   1
Fa0/20    on     802.1q         trunking   1
Fa0/21    on     802.1q         trunking   1
Fa0/22    on     802.1q         trunking   1

Port      Vlans allowed on trunk
Fa0/19    1-99,101-199,201-4094
Fa0/20    1-99,101-199,201-4094
Fa0/21    1-99, 101-299,301-4094
Fa0/22    1-99, 101-299,301-4094
```

Note SW-4 denies VLANs 100 and 200 on ports F0/19-20, whereas, VLANs 100 and 300 are denied on ports F0/21-22.

On SW-1

```
SW-1(config)#int range f0/19-20
SW-1(config-if-range)#Switch trunk allowed vlan except 300
```

To verify the configuration:

On SW-1

SW-1#Sh int trunk

```
Port      Mode   Encapsulation  Status     Native vlan
Fa0/19    on     802.1q         trunking   1
Fa0/20    on     802.1q         trunking   1
Fa0/21    on     802.1q         trunking   1
Fa0/22    on     802.1q         trunking   1

Port      Vlans allowed on trunk
Fa0/19    1-299,301-4094
Fa0/20    1-299,301-4094
Fa0/21    1-199,201-4094
Fa0/22    1-199,201-4094
```

Note VLAN 200 is denied on ports F0/21-22, whereas, VLAN 300 is denied on ports F0/19-20.

The forth Policy:

On SW-1

SW-1(config)#int range F0/19-20
SW-1(config-if-range)#swi trunk allowed vlan except 300,400

SW-1(config)#int range f0/21-22
SW-1(config-if-range)#swi trunk allowed vlan except 200,400

To verify the configuration:

On SW-1

SW-1#Show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/19	on	802.1q	trunking	1
Fa0/20	on	802.1q	trunking	1
Fa0/21	on	802.1q	trunking	1
Fa0/22	on	802.1q	trunking	1

Port	Vlans allowed on trunk
Fa0/19	1-299,301-399,401-4094
Fa0/20	1-299,301-399,401-4094
Fa0/21	1-199,201-399,401-4094
Fa0/22	1-199,201-399,401-4094

Note VLANs 300 and 400 are both denied on ports F0/19-20, whereas, VLANs 200 and 400 are denied on ports F0/21-22.

On SW-2

SW-2(config)#int range f0/19-20
SW-2(config-if-range)#swi trunk allowed vlan except 300,400

To verify the configuration:

On SW-2

SW-2#Sh int trunk

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1
Fa0/21      on           802.1q         trunking      1
Fa0/22      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-299,301-399,401-4094
Fa0/20      1-299,301-399,401-4094
Fa0/21      1-99,101-299,301-4094
Fa0/22      1-99,101-299,301-4094
```

Note VLANs 300 and 400 are denied on ports F0/19-20, whereas, VLANs 100 and 300 are denied on ports F0/21-22.

On SW-3

```
SW-3(config)#int range f0/21-22
SW-3(config-if-range)#swi trunk allowed vlan except 200,400
```

To verify the configuration:

On SW-3

SW-3#Show int trunk

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1
Fa0/21      on           802.1q         trunking      1
Fa0/22      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-99,101-199,201-4094
Fa0/20      1-99,101-199,201-4094
Fa0/21      1-199,201-399,401-4094
Fa0/22      1-199,201-399,401-4094
```

Note VLANs 100 and 200 are both denied on ports F0/19-20, whereas, VLANs 200 and 400 are denied on ports F0/21-22.

Task 8

Configure SW-1 such that it's the root bridge for VLAN 500, if this switch goes down, SW-2 should become the root bridge for this VLAN.

On SW-1

SW-1(config)#Spanning-tree vlan 500 root primary

On SW-2

SW-2(config)#Spanning-tree vlan 500 root secondary

To verify the configuration:

On SW-1

SW-1#Sh spanning-tree vlan 500

```
VLAN0500
  Spanning tree enabled protocol ieee
  Root ID    Priority    25076
             Address     001b.2be5.0e00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    25076  (priority 24576 sys-id-ext 500)
             Address     001b.2be5.0e00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
```

(The rest of the output is omitted)

On SW-2

SW-2#Sh spanning-tree vlan 500

```
VLAN0500
  Spanning tree enabled protocol ieee
  Root ID    Priority    25076
             Address     001b.2be5.0e00
             Cost        19
             Port        21 (FastEthernet0/19)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
  Bridge ID  Priority    29172  (priority 28672 sys-id-ext 500)
             Address     001c.575f.fd00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
```
(The rest of the output is omitted)

Task 9

Ensure that the traffic from SW-2 for VLAN 500 uses ports F0/19 or F0/20 ONLY if the path through SW-4 to SW-3 to SW-1 is NOT possible due to a link being down.

On SW-2

SW-2#Sh spanning-tree vlan 500

```
VLAN0500
  Spanning tree enabled protocol ieee
  Root ID    Priority    25076
             Address     001b.2be5.0e00
             Cost        19
             Port        21 (FastEthernet0/19)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    29172  (priority 28672 sys-id-ext 500)
             Address     001c.575f.fd00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts   Cost      Prio.Nbr  Type
---------------- ---- ----- --------- --------  --------------------------
Fa0/19           Root FWD 19          128.21    P2p
Fa0/20           Altn BLK  19         128.22    P2p
Fa0/21           Desg FWD 19          128.23    P2p
Fa0/22           Desg FWD 19          128.24    P2p
```

Note SW-2 is taking port F0/19 to get to the root bridge for VLAN 500, and it's root cost is 19 which is the cost of a 100 Mbps link. If these ports are shut down, you should see the cost of the local switch (SW-2) to the root bridge, it should be as follows:
The cost of the link from SW-2 to SW-4 which is 19 +

The cost of the link from SW-4 to SW-3 which is 19 +
The cost of the link from SW-3 to SW-1 which is also 19
The total equals to 57, to reveal this information:

On SW-2

```
SW-2(config)#int range f0/19-20
SW-2(config-if-range)#Shut
```

To see the effect:

```
SW-2(config-if-range)#do sh spanning-tree vlan 500

VLAN0500
  Spanning tree enabled protocol ieee
  Root ID    Priority    25076
             Address     001b.2be5.0e00
             Cost        57
             Port        23 (FastEthernet0/21)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    29172  (priority 28672 sys-id-ext 500)
             Address     001c.575f.fd00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 15

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/21           Root FWD 19        128.23   P2p
Fa0/22           Altn BLK 19        128.24   P2p
```

Note it is taking port F0/21 toward SW-4 using port F0/21. If this is traced all the back to SW-1 you will see that the path from SW-2's perspective is through SW-4 to SW-3 to SW-1.

To configure this task, you should "no Shutdown" ports F0/19-20 first.

On SW-2

```
SW-2(config)#int range f0/19-20
SW-2(config-if-range)#No Shut

SW-2(config-if-range)#Spanning-tree vlan 500 cost 58
```

Note the above command sets the cost through ports F0/19-20 higher than 57, and as a result of that, the traffic for VLAN 500 will traverse through the alternate path which is to SW-4 to SW-3 to SW-1 which has a cost of 57.

To verify the configuration:

On SW-2

SW-2#Sh spanning-tree vlan 500

```
VLAN0500
  Spanning tree enabled protocol ieee          Note the best cost is now 57.
  Root ID    Priority    25076
             Address     001b.2be5.0e00
             Cost        57
             Port        23 (FastEthernet0/21)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    29172  (priority 28672 sys-id-ext 500)
             Address     001e.575f.fd00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Altn BLK 58        128.21   P2p
Fa0/20           Altn BLK 58        128.22   P2p
Fa0/21           Root FWD 19        128.23   P2p
Fa0/22           Altn BLK 19        128.24   P2p
```

Task 10

Configure SW-3 as the root bridge for VLAN 600; this switch should be configured such that traffic for VLAN 600 uses the following ports:

F0/21 from SW-1
F0/20 from SW-4

On SW-3

SW-3(config)#int range f0/20-21

```
SW-3(config-if)#Spanning-tree vlan 600 port-priority 0
```

To verify the configuration:

On SW4

```
SW-4#Sh spanning vlan 600

VLAN0600
  Spanning tree enabled protocol ieee
  Root ID    Priority    25176
             Address     000d.65ca.3180
             Cost        19
             Port        20 (FastEthernet0/20)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    33368  (priority 32768 sys-id-ext 600)
             Address     000d.65c1.9200
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Altn BLK 19        128.19   P2p
Fa0/20           Root FWD 19        128.20   P2p
Fa0/21           Altn BLK 19        128.21   P2p
Fa0/22           Altn BLK 19        128.22   P2p
```

On SW-1

```
SW-1#Sh spanning-tree vlan 600

VLAN0600
  Spanning tree enabled protocol ieee
  Root ID    Priority    25176
             Address     000d.65ca.3180
             Cost        19
             Port        23 (FastEthernet0/21)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    33368  (priority 32768 sys-id-ext 600)
             Address     001b.2be5.0e00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
          Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
Fa0/21           Root FWD 19        128.23   P2p
Fa0/22           Altn BLK 19        128.24   P2p
```

Task 11

Erase the config.text and vlan.dat and reload the switches before proceeding to the next lab.

Lab 6 – QinQ Tunneling

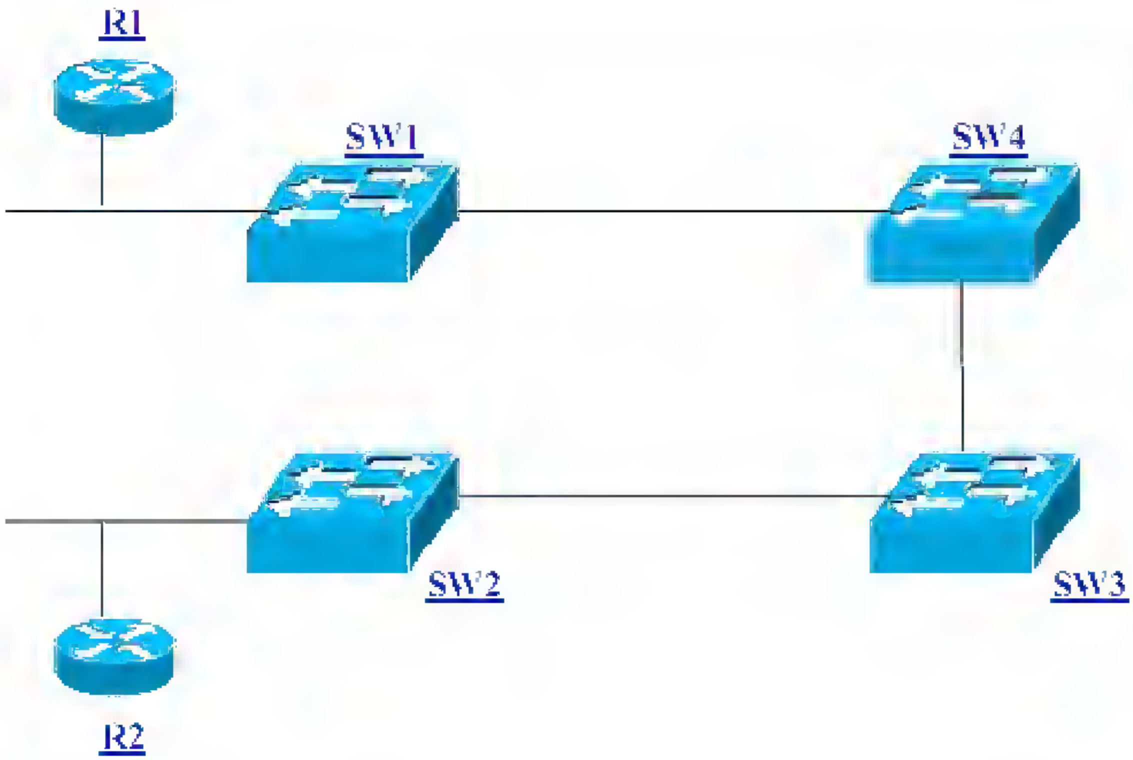

Task 1

Ensure that all interfaces for all 4 switches are in Shutdown mode

On All Switches

```
(config)#int range f0/1-24
(config-if-range)#shutdown
```

Task 2

Establish trunk links between the switches as follows:

1. SW1 and SW4 should use interface F0/23 to establish a trunk link.
2. SW4 and SW3 should use interface F0/19 to establish a trunk link
3. SW3 and SW2 should use interface F0/23 to establish a trunk link

To configure the first item:

On SW1

```
SW1(config)#int f0/23
SW1(config-if)#swi trunk encap isl
SW1(config-if)#swi mode trunk
SW1(config-if)#No shut
```

On SW4

```
SW4(config)#int f0/23
SW4(config-if)#swi trun encap isl
SW4(config-if)#swi mode trunk
SW4(config-if)#No shut
```

To verify the configuration:

On SW1

```
SW1#Show interface trunk

Port      Mode       Encapsulation  Status       Native vlan
Fa0/23    on         isl            trunking     1

Port      Vlans allowed on trunk
Fa0/23    1-4094

Port      Vlans allowed and active in management domain
Fa0/23    1

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/23    1

SW1#Show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme    Capability   Platform   Port ID
SW4              Fas 0/23          125        S I          WS-C3550-2Fas 0/23
```

To configure the second item:

On SW3

```
SW3(config)#int f0/19
SW3(config-if)#swi trunk encap isl
SW3(config-if)#swi mode trunk
SW3(config-if)#no shut
```

On SW4

```
SW4(config)#int f0/19
SW4(config-if)#swi trunk encap isl
SW4(config-if)#swi mode trunk
SW4(config-if)#no shut
```

To verify the configuration:

On SW4

SW4#Show interface trunk

```
Port      Mode      Encapsulation  Status      Native vlan
Fa0/23     on        isl           trunking      1
Fa0/19     on        isl           trunking      1

Port      Vlans allowed on trunk
Fa0/23     1-4094
Fa0/19     1-4094

Port      Vlans allowed and active in management domain
Fa0/23     1
Fa0/19     1

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/23     1
Fa0/19     1
```

SW4#Show cdp neighbor

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
             S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce      Holdtme    Capability    Platform   Port ID
SW1              Fas 0/23            135        S I          WS-C3560-2Fas 0/23
SW3              Fas 0/19            159        S I          WS-C3550-2Fas 0/19
```

To configure the third item:

On SW3

```
SW3(config)#int f0/23
SW3(config-if)#swi trunk encap isl
SW3(config-if)#swi mode trunk
SW3(config-if)#no shut
```

On SW2

```
SW2(config)#int f0/23
SW2(config-if)#swi trunk encap isl
SW2(config-if)#swi mode trunk
SW2(config-if)#no shut
```

To verify the configuration:

On SW3

```
SW3#sh int trunk

Port      Mode      Encapsulation  Status      Native vlan
Fa0/23     on        isl           trunking     1
Fa0/19     on        isl           trunking     1

Port      Vlans allowed on trunk
Fa0/23     1-4094
Fa0/19     1-4094

Port      Vlans allowed and active in management domain
Fa0/23     1
Fa0/19     1

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/23     none
Fa0/19     1

SW3#Sh cdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P – Phone

Device ID        Local Intrfce      Holdtme   Capability   Platform  Port ID
```

SW4	Fas 0/19	143	S I	WS-C3550-2Fas 0/19
SW2	Fas 0/23	165	S I	WS-C3560-2Fas 0/23

On SW1

SW1(config)#VTP domain QinQ

The VTP domain name is configured so that the future VLANs can be propagated across the trunk to the other Switches. You should verify that this configuration is propagated to the other switches before proceeding further. You may need to "Shutdown" and then "No Shutdown" the trunk interface of some of the switches to ensure that they all belong to the same VTP domain.

Task 3

Configure R1's F0/0 and R2's F0/1 using the following IP addresses:
R1's F0/0 = 10.1.12.1 /24 and R2's F0/1 = 10.1.12.2 /24
Ensure that R1 and R2 have full reachability to each other; you should use CDP and Ping to verify.

On R1

R1(config)#int f0/0
R1(config-if)#ip addr 10.1.12.1 255.255.255.0
R1(config-if)#no shut

On R2

R2(config)#int f0/1
R2(config-if)#ip addr 10.1.12.2 255.255.255.0
R2(config-if)#no shut

On SW1

SW1(config)#int f0/1
SW1(config-if)#no shut

On SW2

SW2(config)#int f0/2

SW2(config-if)#no shut

To verify the configuration:

On R1

R1#Sh cdp neighbors

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
SW1              Fas 0/0           154          S I       WS-C3560- Fas 0/1
```

R1#Ping 10.1.12.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

On R2

R2#Show cdp neighbors

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
SW2              Fas 0/1           128          S I       WS-C3560- Fas 0/2
```

R2#Ping 10.1.12.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Note both routers can successfully Ping each other and see the switch that they are directly connected to.

Task 4

Configure these devices such that the output of the "Show cdp neighbor" command on R1 resembles the following:

```
R1#Show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID    Local Intrfce    Holdtme    Capability  Platform    Port ID
R2           Fas 0/0          122        R S I       2611XM      Fas 0/1
```

This task calls for Dot1q Tunneling, 802.1q tunneling enables the service providers to use a single VLAN to support customers who have a single or multiple VLANs that need to connect across the provider's network while preserving their VLAN-IDs. The provider can use this feature to keep traffic from different customers segregated.
When configuring QinQ tunneling, a tunnel port must be defined, this port should be assigned to a VLAN, different customers must be assign to different tunnel ports and different tunnel ports should be configured in different provider VLANs, and this is how the traffic from different customers are segregated.
When a given tunnel port receives customer traffic, it adds a 2 Byte Ether-Type field of 0x8100 followed by a 2 Byte field containing the CoS and the VLAN and this traffic is then put into the VLAN to which the tunnel port is assigned. The Egress tunnel port strips off the 4 Bytes that was added by the ingress tunnel port and transmits the traffic to the customer device.
When Dot1q tunneling is configured, a layer 2 protocol tunneling can also be configured, a layer 2 protocol tunneling allows layer 2 protocol data units (PDUs) to be tunneled through the network, the layer 2 protocols that can be tunneled are: CDP, STP and VTP and they need to be configured or else they will NOT get propagated across the tunnel.
In this case since CDP must be used, it should be configured using the "L2protocol-tunnel CDP" interface configuration command.

On SW1

```
SW1(config-if)#int f0/1
SW1(config-if)#swi acc v 100
SW1(config-if)#swi mode dot1q-tunnel
SW1(config-if)#l2protocol-tunnel cdp
```

On SW2

```
SW2(config-if)#int f0/2
```

```
SW2(config-if)#swi acc v 100
SW2(config-if)#swi mode dot1q-tunnel
SW2(config-if)#l2protocol-tunnel cdp
```

To verify the configuration:

On R1

R1#Show cdp neighbors

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform    Port ID
R2               Fas 0/0           122         R S I       2611XM     Fas 0/1
```

Note sometimes you need to "Clear cdp table" on the routers to see the change.

To verify the tunnel:

On SW1

SW1#Sh dot1q-tunnel

```
dot1q-tunnel mode LAN Port(s)
-----------------------------
Fa0/1
```

SW1#Show l2protocol-tunnel summary

```
COS for Encapsulated Packets: 5
Drop Threshold for Encapsulated Packets: 0

Port    Protocol    Shutdown        Drop           Status
                    Threshold       Threshold
                    (cdp/stp/vtp)   (cdp/stp/vtp)
                    (pagp/lacp/udld) (pagp/lacp/udld)
------- ----------- ---------------- ---------------- -------
Fa0/1    cdp  ---- ---- ----/----/----   ----/----/----   up
         ---- ---- ---- ----/----/----   ----/----/----
```

Note the status is UP and CDP is the ONLY layer 2 protocol tunnel in use.

Task 5

Configure the F0/0 of R1 and F0/1 interface of R2 with two subinterfaces; using the following information:
The first subinterface should belong to VLAN 12, this VLAN on R1's F0/0 interface should have an IP address of 10.1.12.1 /24 and on R2's F0/1 interface it should have an IP address of 10.1.12.2 /24.
The second subinterface should belong to VLAN 34, this VLAN on R1's F0/0 interface should have an IP address of 10.1.34.1 /24 and on R2's F0/1 interface it should have an IP address of 10.1.34.2 /24.
Verify reachability using Ping.

In this task, on each router, two VLANs are created using two subinterfaces, one subinterface is configured to be in VLAN 12 and the other subinterface is configured to be in VLAN 34.

On R1

R1(config)#default interface f0/0

The above command sets the interface f0/0 back to its default configuration.

R1(config)#int f0/0.12
R1(config-subif)#encap dot1q 12
R1(config-subif)#ip addr 10.1.12.1 255.255.255.0

R1(config)#int f0/0.34
R1(config-subif)#encap dot1q 34
R1(config-subif)#ip addr 10.1.34.1 255.255.255.0

On R2

R2(config)#default interface f0/1

The above command sets the interface f0/0 back to its default configuration.

R2(config)#int f0/1.12
R2(config-subif)#encap dot1q 12
R2(config-subif)#ip addr 10.1.12.2 255.255.255.0

R2(config)#int f0/1.34
R2(config-subif)#encap dot1q 34
R2(config-subif)#ip addr 10.1.34.2 255.255.255.0

Because the interfaces of the routers are configured as a trunk ports, the switchport that they connect to is configured to be in VLAN 100, and they are also configured as a tunnel ports, this is called an asymmetrical link.

To verify the configuration:

On R1

R1#Ping 10.1.12.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (4/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 10.1.34.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (4/5), round-trip min/avg/max = 1/2/4 ms

Task 6

Delete the VLAN.dat and Config.text of SW1 and SW2 and reload these switches, while the switches are reloading configure the following:
Set the f0/23 interface on SW3 and SW4 to default configuration and authenticate the VTP domain using "QinQ" as the password.
Ensure that ONLY ports F0/23 and F0/1 on SW1 and F0/23 and F0/2 on SW2 are in UP/UP state, the rest of the ports should be disabled.

On SW1 and SW2

#delete config.text
#delete vlan.dat
#reload

On SW1

Switch(config)#host SW1

On SW2

```
Switch(config)#host SW2
```

On SW3 and SW4

```
(config)#default interface f0/23
(config)#VTP password QinQ
```

On SW1

```
SW1(config)#interface range f0/2-24
SW1(config-if-range)#Shutdown
```

On SW2

```
SW2(config)#int range f0/1, f0/3-24
SW2(config-if-range)#Shutdown
```

Task 7

Configure these devices such that the output of the "Show cdp neighbor" command on SW1 & SW2 resembles the following:

On SW1

```
SW1#Show cdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme   Capability   Platform   Port ID
SW2              Fas 0/23           172         S I        WS-C3560-2Fas 0/23
R1               Fas 0/1            144        R S I       2611XM    Fas 0/0
```

On SW2

```
SW2#Show cdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
```

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
SW1	Fas 0/23	174	S I	WS-C3560-2	Fas 0/23
R2	Fas 0/2	161	R S I	2611XM	Fas 0/1

On SW3 and SW4

```
(config)#int f0/23
(config-if)#swi acc v 100
(config-if)#swi mode dot1q-tunnel
(config-if)#l2protocol-tunnel cdp
```

To verify the configuration:

On SW1

SW1#Show cdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
SW2	Fas 0/23	172	S I	WS-C3560-2	Fas 0/23
R1	Fas 0/1	144	R S I	2611XM	Fas 0/0

On SW2

SW2#Show cdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
SW1	Fas 0/23	174	S I	WS-C3560-2	Fas 0/23
R2	Fas 0/2	161	R S I	2611XM	Fas 0/1

Task 8

Configure a trunk link between SW1 and SW2 using interface F0/23. These two switches should use Dot1Q encapsulation for this task.

On SW1 and SW2

```
(config)#int f0/23
(config-if)#swi trunk encap dot1q
(config-if)#swi mode trunk
```

To verify the configuration:

On SW2

SW2#Show interface trunk

```
Port     Mode      Encapsulation  Status      Native vlan
Fa0/23   on        802.1q         trunking     1

Port     Vlans allowed on trunk
Fa0/23    1-4094

Port     Vlans allowed and active in management domain
Fa0/23     1

Port     Vlans in spanning tree forwarding state and not pruned
Fa0/23     1
```

Task 9

Configure SW1 and SW2 in VTP domain called "Customer"; Configure VLANs 35 and 46 on SW1 and ensure that they are propagated to SW2. SW1 should be the root bridge for 35 and SW2 should be the root bridge for 46.

To accomplish this task, VTP and STP protocols should be propagated between SW1 and SW2; this is accomplished by configuring SW3 and SW4 to allow these protocols on their tunnel port (F0/23) as follows:

On SW3 and SW4

```
(config)#inter f0/23
(config-if)#L2protocol-tunnel STP
(config-if)#L2protocol-tunnel VTP
```

On SW1

```
SW1(config)#VTP domain Customer
```

To verify the configuration:

On SW2

```
SW2#Sh vtp status

VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 5
VTP Operating Mode              : Server
VTP Domain Name                 : Customer
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00
Local updater ID is 0.0.0.0 (no valid interface found)
```

To create VLANs 35 and 46:

On SW1

```
SW1(config)#VLAN 35,46
SW1(config)#exit
```

To verify the configuration:

On SW2

```
SW2#Sh vlan br | exc unsup

VLAN Name                             Status    Ports

---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                                Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                                Fa0/14, Fa0/15, Fa0/16, Fa0/17
```

Fa0/18, Fa0/19, Fa0/20, Fa0/21
Fa0/22, Fa0/23, Fa0/24, Gi0/1
Gi0/2

```
35   VLAN0035                  active
46   VLAN0046                  active
```

To determine the root bridge for the newly created VLANs:

On SW1

SW1#Sh spanning-tree vlan 35

```
VLAN0035
  Spanning tree enabled protocol ieee
  Root ID    Priority    32803
             Address     001a.2f0a.2000
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32803  (priority 32768 sys-id-ext 35)
             Address     001a.2f0a.2000
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/7            Desg FWD 19        128.9    P2p
```

SW1#Sh spanning-tree vlan 46

```
VLAN0046
  Spanning tree enabled protocol ieee
  Root ID    Priority    32814
             Address     001a.2f0a.2000
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32814  (priority 32768 sys-id-ext 46)
             Address     001a.2f0a.2000
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
```

```
Fa0/7            Desg  FWD 19       128.9      P2p
```

Note in this case SW2 is the root bridge for both (35 and 46) VLANs. The result may be different on your pod of routers and switches.

The last step is to configure SW1 as the root bridge for VLAN 35 and SW2 as the root bridge for VLAN 46, as follows:

On SW1

```
SW1(config)#Spanning-tree vlan 35 root primary
```

To verify the configuration:

On SW2

```
SW2#Sh spanning-tree vlan 35

VLAN0035
  Spanning tree enabled protocol ieee
  Root ID    Priority    24611
             Address     0023.050b.e780
             Cost        19
             Port        9 (FastEthernet0/7)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32803  (priority 32768 sys-id-ext 35)
             Address     001a.2f0a.2000
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role  Sts   Cost      Prio.Nbr Type
---------------- ----  ---- --------- -------- --------------------------------
Fa0/23           Desg  FWD 19        128.9      P2p
```

On SW1

```
SW1#Sh spanning-tree vlan 35

VLAN0035
  Spanning tree enabled protocol ieee
  Root ID    Priority    24611
             Address     0023.050b.e780
```

```
        This bridge is the root
        Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24611  (priority 24576 sys-id-ext 35)
        Address     0023.050b.c780
        Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
        Aging Time 300                         Note SW1 is the root bridge

Interface        Role  Sts   Cost      Prio.Nbr Type
---------------- ------ ------ --------- -------- --------------------------------
Fa0/7            Desg   FWD 19        128.9      P2p
```

On SW2

SW2(config)#Spanning-tree vlan 46 root primary

To verify the configuration:

On SW2

SW2#Sh spanning vlan 46

```
VLAN0046
  Spanning tree enabled protocol ieee
  Root ID    Priority    24622            Note this switch is the root bridge for VLAN 46
        Address     001a.2f0a.2000
        This bridge is the root
        Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24622  (priority 24576 sys-id-ext 46)
        Address     001a.2f0a.2000
        Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
        Aging Time 300

Interface        Role  Sts   Cost      Prio.Nbr Type
---------------- ------ ------ --------- -------- --------------------------------
Fa0/7            Desg   FWD 19        128.9      P2p
```

Task 10

Delete VLAN.dat and config.text on all switches and reload them before proceeding to the next lab.

Lab 7 – Fallback Bridging

Lab Setup:

- Use the chart below for address assignment:

IP Addressing:

Router	Interface	IPX Net address	IPv6 Address	VLAN	Mac-address
BB2	FastEthernet	ABCD	23::2 /64	20	0000.2222.2222
BB3	FastEthernet	ABCD	23::3 /64	30	0000.3333.3333

Task 1

Shutdown all the used ports on SW1 and SW3, only ports that are connected to BB2 and BB3 should be in UP/UP state.

On SW1

```
SW1(config)#int range f0/1 - 9 , f0/12 – 24
SW1(config-if-range)#Shut
```

On SW3

```
SW1(config)#int range f0/1 - 11 , f0/14 – 24
```

SW1(config-if-range)#Shut

Task 2

Configure the appropriate switch such that routers BB2 and BB3 can forward NON-IP traffic between VLAN 20 and 30; Fallback Bridging should be configured to accomplish this task. If this task is configured properly, you should be able to use "Ping" to test this configuration using IPv6 or IPX addressing identified in the IP addressing chart.

Note since the task specifies that the test should be conducted using IPv6 and IPX, 3550 switches will be the only choice. Since the earlier IOS versions did NOT have support for IPv6, these switches looked at IPv6 traffic as NON-IP, just like IPX.

To configure Fallback Bridging:

On SW3

The following command assigns a bridge group number (In this case number 1) and it also specifies the VLAN bridge spanning-tree protocol to run in this bridge group.

SW3(config)#**bridge 1 protocol vlan-bridge**

The following configuration assigns the bridge group that was created with the "Bridge 1 protocol vlan-bridge" global configuration command to interface VLAN 20 and 30.

SW3(config)#int vlan 20
SW3(config-if)#**bridge-group 1**

SW3(config-if)#int vlan 30
SW3(config-if)#**bridge-group 1**

To verify the configuration

On SW3

If the output of your "Show bridge" command does NOT reveal the MAC address of BB2 and BB3, you should generate some traffic (For example: Pinging BB3 from BB2 using the IPv6 or IPX) so the bridge will see the MAC addresses.

SW3#Show bridge

```
Br Group    Mac Address       State     Type        Ports
--------    --------------    -------   ------      ------
   1        0000.2222.2222    Forward   DYNAMIC     Vl20 Fa0/12
   1        0000.3333.3333    Forward   DYNAMIC     Vl30 Fa0/13
```

To test the configuration:

On BB2

BB2#Ping 23::3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

BB2#Ping IPX ABCD.0000.3333.3333

```
Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to ABCD.0000.3333.3333, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 3

Configure the switch such that ONLY static entries are bridged, if this switch is configured properly, the switch should NOT bridge dynamically learnt Mac addresses.

On SW3

In the previous task, the switch (SW3) learned the MAC addresses dynamically, and it bridged the traffic between the VLANs. The following command prevents the switch to forward frames to stations that it has learned dynamically.

SW3(config)#**NO bridge 1 acquire**

To verify the configuration:

Note the output of the following "Show" command reveals that the dynamically learned MAC addresses are discarded:

On SW3

SW3#Show bridge

Br Group	Mac Address	State	Type	Ports
--------	----------------	-------	------	------
1	0000.2222.2222	discard	DYNAMIC	Vl20 Fa0/12
1	0000.3333.3333	discard	DYNAMIC	Vl30 Fa0/13

To test the configuration:

On BB2

BB2#Ping IPX ABCD.0000.3333.3333

Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to ABCD.0000.3333.3333, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB2#Ping 23::3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23::3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

To complete the configuration:

The following two commands add the MAC addresses of BB2 and BB3

statically, therefore, since the traffic from dynamically learned MAC addresses are discarded, the traffic with statically configured MAC addresses will be forwarded.

On SW3

SW3(config)#Bridge 1 address 0000.2222.2222 forward
SW3(config)#Bridge 1 address 0000.3333.3333 forward

To verify the configuration:

On BB2

```
SW3#Show bridge

Br Group   Mac Address      State     Type     Ports
--------   --------------   -------   ------   ------
   1       0000.2222.2222   Forward   Static   -
   1       0000.3333.3333   Forward   Static   -
```

To test the configuration:

```
BB2#Ping 23::3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/1/4 ms

BB2#Ping IPX ABCD.0000.3333.3333

Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to ABCD.0000.3333.3333, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 4

Configure the appropriate switch such that routers BB2 and BB3 can forward NON-IP traffic between VLAN 20 and 30; you should configure Fallback Bridging to accomplish this task. If this task is configured properly, you should be able to use "Ping" to test this configuration using IPX addressing identified in the addressing chart. IPv6 addressing should NOT work when conducting tests using the Ping command.

Note because 3560 switches support IPv6, they do not consider IPv6 as NON-IP traffic; therefore, they do not bridge IPv6 traffic.

On BB2

```
BB2(config)#default interface f0/1

BB2(config)#int f0/0
BB2(config-if)#mac-address 000.2222.2222
BB2(config-if)#ipx Network ABCD
```

```
BB2(config-if)#ipv6 address 23::2/64
BB2(config-if)#no shut
```

On BB3

```
BB3(config)#default interface f0/1

BB3(config)#int f0/0
BB3(config-if)#mac-address 0000.3333.3333
BB3(config-if)#ipx Network ABCD
BB3(config-if)#ipv6 address 23::3/64
BB3(config-if)#no shut
```

On SW1

```
SW1(config)#int f0/11
SW1(config-if)#swi mode acc
SW1(config-if)#swi acc v 20

SW1(config-if)#int f0/12
SW1(config-if)#swi mode acc
SW1(config-if)#swi acc v 30

SW1(config)#int vlan 20
SW1(config-if)#bridge-group 1

SW1(config-if)#int vlan 30
SW1(config-if)#bridge-group 1

SW1(config)#Bridge 1 protocol vlan-bridge
```

To verify the configuration:

On SW1

SW3#Show bridge

Br Group	Mac Address	State	Type	Ports
1	0000.2222.2222	Forward	DYNAMIC	Vl20
1	0000.3333.3333	Forward	DYNAMIC	Vl30

To test the configuration:

On SW1

BB2#Ping 23::3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23::3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note the above Ping failed but the following Ping worked.

BB2#Ping ipx ABCD.0000.3333.3333

Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to ABCD.0000.3333.3333, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Task 5

Configure R1 based on the following chart:

Router	Interface	IPX Net address	IPv6 Address	VLAN	Mac-address
R1	Fast Ethernet	ABCD	23::1 /64	Default	0000.1111.1111

On R1

R1(config)#ipx routing

R1(config)#int f0/0
R1(config-if)#mac-address
R1(config-if)#mac-address 0000.1111.1111
R1(config-if)#ipx Network ABCD
R1(config-if)#ipv6 address 23::1/64
R1(config-if)#no shut

On SW1

SW1(config)#interface f0/0
SW1(config-if)#no Shut

```
SW1(config)#int vlan 1
SW1(config-if)#bridge-group 1
SW1(config-if)#no shut
```

To test the configuration:

On R1

```
R1#ping ipx abcd.0000.2222.2222

Type escape sequence to abort.
Sending 5, 100-byte IPX Novell Echoes to ABCD.0000.2222.2222, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

To verify the configuration:

On SW1

```
SW1#Show bridge


Br Group  Mac Address       State     Type        Ports
--------  ----------------  -------   ------      ------
    1     0000.1111.1111    Forward   DYNAMIC     Vl1
    1     0000.2222.2222    Forward   DYNAMIC     Vl20
    1     0000.3333.3333    Forward   DYNAMIC     Vl30
```

Task 6

Erase the startup configuration on the routers, Switches and reload them before proceeding to the next task.

Lab 8
Multiple Spanning Trees (802.1s)

Task 1

The first Catalyst switch should be configured with a hostname of Cat-1 and the second Catalyst should have a hostname of Cat-2.

On the first Switch

Switch(config)#Hostname Cat-1

On the Second Switch

Switch(config)#Hostname Cat-2

Task 2

Configure ports F0/21-24 on Cat-1 and Cat-2 in shutdown state.

On Both Switches

Cat-2(config)#int range F0/21-24
Cat-2(config-if-range)#Shut

Task 3

Ports F0/19-20 on both switches should be in trunking mode, these ports should use an industry standard protocol to establish the trunk.

On Both Switches:

(config)#int range f0/19-20

```
(config-if-range)#Switchport trunk encapsulation dot1q
Cat-1(config-if-range)#Switchport mode trunk
```

Task 4

Create VLANs 12, 34, 56, and 90 on Cat-1 and ensure that these VLANs are propagated to Cat-2 via VTP messages.

On Cat-1

```
Cat-1(config)#vlan 12,34,56,90
Cat-1(config-vlan)#exit
```

To verify the configuration:

On Cat-1

```
Cat-1#Sh vlan br | exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12, Fa0/13, Fa0/14
                                                Fa0/15, Fa0/16, Fa0/17, Fa0/18
                                                Fa0/19, Fa0/20, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
12   VLAN0012                         active
34   VLAN0034                         active
56   VLAN0056                         active
90   VLAN0090                         active
```

On Cat-2

```
Cat-2#Sh vlan br | exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/9, Fa0/10
```

```
                                                Fa0/11, Fa0/12, Fa0/13, Fa0/14
                                                Fa0/15, Fa0/16, Fa0/17, Fa0/18
                                                Fa0/19, Fa0/20, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
```

Note none of the switches are in VTP transparent mode, and yet the VLANs are not getting propagated from Cat-1 to Cat-2. This is because the VTP domain name is not configured, if the VTP domain name is not configured, the switches will NOT propagate their VLAN information across the trunk links. For the purpose of this lab VTP domain name of "TST" is created so Cat-1 propagates the VLAN information to Cat-2.

On Cat-1

```
Cat-1(config)#vtp domain TST
```

Note the above command configures a VTP domain name, if the other switch does not have a domain name configured and a trunk has been established between the two switches, Cat-1 will convey the domain name via VTP messages and the two switches will synch up their VLAN information based on the highest VTP rev number. In this task, since a name has not been specified, a domain name of "TST" has been configured.

To verify the configuration:

On Cat-2

```
Cat-2#Sh vlan brief | exc unsup


VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12, Fa0/13, Fa0/14
                                                Fa0/15, Fa0/16, Fa0/17, Fa0/18
                                                Fa0/19, Fa0/20, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
12   VLAN0012                         active
34   VLAN0034                         active
56   VLAN0056                         active
90   VLAN0090                         active
```

Task 5

Configure all the ports except F0/19 and F0/20 in shutdown mode.

On Both Switches

```
Cat-x(config)#int range f0/1-18 , F0/21-24
Cat-x(config-if-range)#Shut
```

Task 6

Configure Multi-instance of Spanning Tree on these two switches using the follows policy:

1. There should be two instances of STP, instance 1 and 2
2. The revision number should be 1
3. The MST region name should be "CCIE"
4. Instance 1 should handle VLANs 12 and 34
5. Instance 2 should handle VLAN 56
6. All future VLANs should use instance 0
7. Instance 1 should use F0/19
8. Instance 2 should use F0/20
9. Cat-1 should be the root bridge for the first instance
10. Cat-2 should be the root bridge for the second instance

On Both Switches

The default mode for spanning-tree is PVST, the output of the following Show command verifies this information:

#Show spanning-tree summary

The default mode of Spanning-tree

Switch is in pvst mode
Root bridge for: none
Extended system ID is enabled
(The rest of the output is omitted)

On Both Switches

(config)#**Spanning-tree mode mst**

This command enables and changes the mode of the spanning-tree on the switch to MST.

To verify the configuration:

On Both Switches:

#Sh spanning-tree sum

Switch is in mst mode (IEEE Standard)
(The rest of the output is omitted)

To configure MST on the switches:

On Both Switches:

(config)#Spanning-tree mst configuration

The above command enters the MST configuration mode

(config-mst)#Revision 1

The above command sets the MST configuration revision number to 1. The range for this number is 1-65535.

(config-mst)#Name CCIE

The above command configured the name of the region to be "CCIE"

(config-mst)#Instance 1 vlan 12,34
(config-mst)#Instance 2 vlan 56
(config-mst)#exit

MST supports 16 instances, once the spanning-tree mode is changed to MST and the MST configuration mode is entered, instance 0 is created and all VLANs are mapped to that instance. The above commands map the requested VLAN/s to the specified instances, and by default all the future VLANs or VLAN/s that are not statically mapped to a given instance will be assigned to instance 0, instance 0 is the Catch all instance.

To verify this configuration:

On both Switches

```
#Show spanning-tree mst configuration

Name      [CCIE]
Revision  1              Instances configured 3
Instance  Vlans mapped
--------  ---------------------------------------
0         1-11,13-33,35-55, 57-4094
1         12,34
2         56
```

To Verify the configuration before configuring the next portion of the task:

On Cat-1

```
Cat-1#Show spanning-tree bridge
                                                   Hello  Max  Fwd
MST Instance          Bridge ID                    Time   Age  Dly   Protocol
---------------- --------------------------------- -----  ---- ---   ----------
MST0             32768 (32768,  0) 0015.639d.5880    2     20   15    mstp
MST1             32769 (32768,  1) 0015.639d.5880    2     20   15    mstp
MST2             32770 (32768,  2) 0015.639d.5880    2     20   15    mstp
```

Note this command displays the BID for your switch (This is NOT the BID of the root Bridge), and instead of assigning a BID to each VLAN, there is a BID for each instance, the priority is incremented based on the instance number, this is the only time that we see a priority value of 32768 assigned to a VLAN or a group of VLANs.

To see the root bridge for a given instance:

On Cat-1

```
Cat-1#Show spanning-tree root
                                          Root      Hello  Max   Fwd
MST Instance          Root ID             Cost      Time   Age   Dly  Root Port
---------------- -------------------- ---------- -------- ------ ---- ------------
MST0             32768  0015.639d.5880         0      2     20    15
MST1             32769  0015.639d.5880         0      2     20    15
MST2             32770  0015.639d.5880         0      2     20    15
```

On Cat-2

Cat-2#Show spanning-tree root

```
                                          Root      Hello  Max   Fwd
MST Instance        Root ID               Cost      Time   Age   Dly  Root Port
---------------- ------------------------ --------- ------ ----- ---- ------------
MST0           32768  0015.639d.5880           0      2     20    15   F0/19
MST1           32769  0015.639d.5880      200000      2     20    15   F0/19
MST2           32770  0015.639d.5880      200000      2     20    15   F0/19
```

The above command displays the BID of the root bridge for different instances. The output may vary based on the Switch's BID.

Enter the following command to see which switch has a BID value of "0015.639d.5880":

Cat-1#Sh version | Inc Base ethernet

Base ethernet MAC Address : 0015:639D:5880

On Cat-1

Cat-1(config)#Spanning-tree mst 1 priority 0
Cat-1(config)#Spanning-tree mst 2 priority 4096

On Cat-2

Cat-2(config)#Spanning-tree mst 1 priority 4096
Cat-2(config)#Spanning-tree mst 2 priority 0

The above commands will change the switch priority such that Cat-1 will be chosen as the root switch for instance 1 and Cat-2 will be chosen as the root bridge for instance 2.
By default the "Spanning-tree extend system-id" is configured as part of your startup configuration, because the extended system id is set, the priority must be configured in increments of 4096. Remember the lower value has higher preference.

To verify the configuration:

On Cat-1

Cat-1#Show spanning root

```
                                          Root      Hello Max Fwd
MST Instance          Root ID             Cost      Time  Age Dly  Root Port
---------------- ------------------------ --------- ----- ------- --- ---------
```

```
MST0         32768  0015.639d.5880         0     2   20   15
MST1             1  0015.639d.5880         0     2   20   15
MST2             2  001c.f901.3d80    200000     2   20   15  Fa0/19
```

The local switch (Cat-1) is the root bridge for instance 0 and 1 (This may be different based on your switch's BID). The column that specifies the Root ID shows the priority for MST1 and MST2 as 1 and 2 respectively, the priority is the sum of instance number plus the Priority. Remember that this switch's priority is set to zero.
Note this switch is not the root for MST2. Another indication that it is not the root for instance 2 is the root port, remember that the root bridge does not have any ports set as root for the VLANs or in this case Instances that it's the root bridge for.

On Cat-2

```
Cat-2#Show spanning root

                                           Root    Hello Max Fwd
MST Instance          Root ID              Cost    Time  Age Dly  Root Port
---------------- -------------------- --------- ----- ------- --- ----------
MST0         32768  0015.639d.5880         0     2   20   15  Fa0/19
MST1             1  0015.639d.5880    200000     2   20   15  Fa0/19
MST2             2  001c.f901.3d80         0     2   20   15
```

Note Cat-2 is the root bridge for instance 2, whereas, Cat-1 is the root for MST instances of 0 and 1.

To configure the last portion of this task, the existing state is displayed in the output of the following show command:

On Cat-1

```
Cat-1#Show spanning int f0/19

Mst Instance     Role  Sts    Cost      Prio.Nbr Type
---------------- ------- ------- --------- -------- -------
MST0             Desg  FWD  200000    128.21   P2p
MST1             Desg  FWD  200000    128.21   P2p
MST2             Root  FWD  200000    128.21   P2p

Cat-1#Show spanning int f0/20

Mst Instance     Role  Sts    Cost      Prio.Nbr Type
---------------- ------- ------- --------- -------- -------
```

```
MST0         Desg   FWD  200000    128.22   P2p
MST1         Desg   FWD  200000    128.22   P2p
MST2         Altn   BLK  200000    128.22   P2p
```

On Cat-2

Cat-2#Show spanning int f0/19

```
Mst Instance    Role  Sts    Cost      Prio.Nbr Type
--------------- ------- ------- --------- --------- -------
MST0         Root   FWD  200000    128.21   P2p
MST1         Root   FWD  200000    128.21   P2p
MST2         Desg   FWD  200000    128.21   P2p
```

Cat-2#Show spanning int f0/20

```
Mst Instance    Role  Sts    Cost      Prio.Nbr  Type
--------------- ------- ------- --------- ---------  -------
MST0         Altn   BLK  200000    128.22   P2p
MST1         Altn   BLK  200000    128.22   P2p
MST2         Desg   FWD  200000    128.22   P2p
```

Note based on the output of the above show commands, traffic for all MST instances take port F0/19 and none of the instances are using port f0/20.
To configure items 7 and 8, port-priority command is used as follows:

On Both switches

```
(config)#Int F0/19

(config-if)#Spanning-tree mst 1 port-priority 0
(config-if)# Spanning-tree mst 2 port-priority 128

(config)#Int F0/20
(config-if)#Spanning-tree mst 1 port-priority 128
(config-if)#Spanning-tree mst 2 port-priority 0
```

High priority

In this task Port-priority is used when selecting an interface to put into the forwarding state for a given instance; a lower value has a higher priority.
In this case port F0/19 will be used by all the VLANs that are assigned to instances 0 & 1, because it has a higher priority (Lower value), and instance 2 will use port F0/20 because it has been configured with a higher priority (Lower value).

To verify the configuration:

On Cat-1

Cat-1#Show spanning-tree int f0/19

Mst Instance	Role	Sts	Cost	Prio.Nbr	Type
MST0	Desg	FWD	200000	128.21	P2p
MST1	Desg	FWD	200000	0.21	P2p
MST2	Altn	BLK	200000	128.21	P2p

Cat-1#Sh spanning-tree int f0/20

Mst Instance	Role	Sts	Cost	Prio.Nbr	Type
MST0	Desg	FWD	200000	128.22	P2p
MST1	Desg	FWD	200000	128.22	P2p
MST2	Root	FWD	200000	0.22	P2p

On Cat-2

Cat-2#Show spanning-tree int f0/19

Mst Instance	Role	Sts	Cost	Prio.Nbr	Type
MST0	Root	FWD	200000	128.21	P2p
MST1	Root	FWD	200000	0.21	P2p
MST2	Desg	FWD	200000	128.21	P2p

Cat-2#Sh spanning-tree int f0/20

Mst Instance	Role	Sts	Cost	Prio.Nbr	Type
MST0	Altn	BLK	200000	128.22	P2p
MST1	Altn	BLK	200000	128.22	P2p
MST2	Desg	FWD	200000	0.22	P2p

Note instances 0 & 1 use port F0/19 whereas, instance 2 uses port f0/20.

Task 6

Erase the startup configuration and vlan.dat before proceeding to the next lab

Lab 9
Private VLANs

Task 1

The first switch should be configured with a hostname of SW1 and the second switch should be configured with a hostname of SW2

On the First Switch

Switch(config)#Hostname SW1

On the Second Switch

Switch(config)#Hostname SW2

Task 2

Shutdown ports F0/21-24 on SW1 and SW2

On Both Switches:

(config)#int range f0/21-24
(config-if-range)#shut

Task 3

Configure trunking between SW1 and SW2 using ports F0/19 and F0/20. Use an industry standard trunking protocol for this purpose. Assign a brief meaningful description to these interfaces.

On Both Switches

SWx(config)#Interface range f0/19-20

```
SWx(config-if-range)#Switch trunk encap dot1q
SWx(config-if-range)#Switch mode trunk
SWx(config-if-range)#Description Trunk to SWx
```

Note you should replace the "x" on "SWx" in the description with the appropriate Switch number.

----------------------------Recommendation-----------------------
If the description is configured for each interface, the output of the "Show interface status" can help understand the topology of the lab.

To verify the configuration:

On SW1

SW1#Show int trunk

```
Port      Mode         Encapsulation  Status        Native vlan
Fa0/19    on           802.1q         trunking      1
Fa0/20    on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/20      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1
Fa0/20      1

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
Fa0/20      none
```

On SW2

SW2#Show int trunk

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/20      1-4094
```

```
Port        Vlans allowed and active in management domain
Fa0/19      1
Fa0/20      1

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
Fa0/20      1
```

Task 4

Assign IP addressing to the interface of the routers using the following chart and ensure that these routers can ping each other. You should assign a brief meaningful interface description on the switchports.

Router	Interface	IP address and Subnet mask
R1	F0/0	200.1.1.1 /24
R2	F0/0	200.1.1.2 /24
R3	F0/1	200.1.1.3 /24
R4	F0/0	200.1.1.4 /24
R5	F0/1	200.1.1.5 /24
R6	F0/1	200.1.1.6 /24
BB1	F0/1	200.1.1.7 /24
BB2	F0/0	200.1.1.8 /24
BB3	F0/0	200.1.1.9 /24

On R1

```
R1(config)#Int F0/0
R1(config-if)#Ip address 200.1.1.1 255.255.255.0
R1(config-if)#No shut
```

On R2

```
R2(config)#Int F0/0
R2(config-if)#Ip address 200.1.1.2 255.255.255.0
R2(config-if)#No shut
```

On R3

```
R3(config)#Int F0/1
```

R3(config-if)#Ip address 200.1.1.3 255.255.255.0
R3(config-if)#No shut

On R4

R4(config)#Int F0/0
R4(config-if)#Ip address 200.1.1.4 255.255.255.0
R4(config-if)#No shut

On R5

R5(config)#Int F0/1
R5(config-if)#Ip address 200.1.1.5 255.255.255.0
R5(config-if)#No shut

On R6

R6(config)#Int F0/1
R6(config-if)# Ip address 200.1.1.6 255.255.255.0
R6(config-if)# No shut

On BB1

BB1(config)#Int F0/1
BB1(config-if)# Ip address 200.1.1.7 255.255.255.0
BB1(config-if)# No shut

On BB2

BB2(config)#int F0/0
BB2(config-if)#ip address 200.1.1.8 255.255.255.0
BB2(config-if)#No shut

On BB3

BB3(config)#int F0/0
BB3(config-if)#ip address 200.1.1.9 255.255.255.0
BB3(config-if)#No shut

On SW1

SW1(config)#Int F0/1

SW1(config-if)#Description R1's F0/0

SW1(config)#Int F0/2
SW1(config-if)#Description R2's F0/0

SW1(config)#Int range F0/3 , F0/5-9 , F0/12-18 , F0/21-24
SW1(config-if-range)#Description --

SW1(config)#Int F0/4
SW1(config-if)#Description R4's F0/0

SW1(config)#Int F0/12
SW1(config-if)#Description BB2's F0/0

SW1(config)#Int F0/13
SW1(config-if)#Description BB3's F0/0

SW1(config)#Int range F0/19-20
SW1(config-if-range)#Description Trunk to SW2

On SW2

SW2(config)#Int range F0/1-2 , F0/4 , F0/10-18 , F0/21-24
SW2(config-if-range)#Description --

SW2(config)#Int F0/3
SW2(config-if)#Description R3's F0/1

SW2(config)#Int F0/5
SW2(config-if)#Description R5's F0/1

SW2(config)#Int F0/6
SW2(config-if)#Description R6's F0/1

SW2(config)#Int F0/11
SW2(config-if)#Description BB1's F0/1

SW2(config)#Int range F0/19-20
SW2(config-if-range)#Description Trunk to SW1

To test and verify the configuration:

On R1

```
R1#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.8

Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 5

Configure the switches such that the ports that are not used are in administratively down state. Use minimum number of commands for this task.

On SW1

```
SW1(config)#int range F0/3 , F0/5 , F0/10, F0/14-18 , F0/21-24
SW1(config-if-range)#Shut
```

To verify the configuration:

On SW1

SW1#Sh int status

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/1	R1's F0/0	connected	1	a-full	a-100	10/100BaseTX
Fa0/2	R2's F0/0	connected	1	a-full	a-100	10/100BaseTX
Fa0/3	--	disabled	1	auto	auto	10/100BaseTX
Fa0/4	R4's F0/0	connected	1	a-full	a-100	10/100BaseTX
Fa0/5	--	disabled	1	auto	auto	10/100BaseTX
Fa0/6	--	disabled	1	auto	auto	10/100BaseTX
Fa0/7	--	disabled	1	auto	auto	10/100BaseTX
Fa0/8	--	disabled	1	auto	auto	10/100BaseTX
Fa0/9	--	disabled	1	auto	auto	10/100BaseTX
Fa0/10	--	disabled	1	a-full	a-100	10/100BaseTX
Fa0/11	BB1's F0/0	connected	1	a-full	a-100	10/100BaseTX

Fa0/12	BB2's F0/0	disabled	1	auto	auto	10/100BaseTX
Fa0/13	BB3's F0/0	disabled	1	auto	auto	10/100BaseTX
Fa0/14	--	disabled	1	auto	auto	10/100BaseTX
Fa0/15	--	disabled	1	auto	auto	10/100BaseTX
Fa0/16	--	disabled	1	auto	auto	10/100BaseTX
Fa0/17	--	disabled	1	auto	auto	10/100BaseTX
Fa0/18	--	disabled	1	auto	auto	10/100BaseTX
Fa0/19	Trunk to SW2	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/20	Trunk to SW2	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/21	--	disabled	1	auto	auto	10/100BaseTX
Fa0/22	--	disabled	1	auto	auto	10/100BaseTX
Fa0/23	--	disabled	1	auto	auto	10/100BaseTX
Fa0/24	--	disabled	1	auto	auto	10/100BaseTX

On SW2

SW2(config)#int range F0/1-2 , F0/4 , F0/8-10, F0/12-18 , F0/21-24
SW2(config-if)#Shut

To verify the configuration:

On SW2

SW2#Sh int status

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/1	--	connected	1	a-full	a-100	10/100BaseTX
Fa0/2	--	connected	1	a-full	a-100	10/100BaseTX
Fa0/3	R3's F0/1	disabled	1	auto	auto	10/100BaseTX
Fa0/4	--	connected	1	a-full	a-100	10/100BaseTX
Fa0/5	R5's F0/1	disabled	1	auto	auto	10/100BaseTX
Fa0/6	R6's F0/1	disabled	1	auto	auto	10/100BaseTX
Fa0/7	--	disabled	1	auto	auto	10/100BaseTX
Fa0/8	--	disabled	1	auto	auto	10/100BaseTX
Fa0/9	--	disabled	1	auto	auto	10/100BaseTX
Fa0/10	--	connected	1	a-full	a-100	10/100BaseTX
Fa0/11	BB1's F0/0	connected	1	a-full	a-100	10/100BaseTX
Fa0/12	--	disabled	1	auto	auto	10/100BaseTX
Fa0/13	--	disabled	1	auto	auto	10/100BaseTX
Fa0/14	--	disabled	1	auto	auto	10/100BaseTX
Fa0/15	--	disabled	1	auto	auto	10/100BaseTX
Fa0/16	--	disabled	1	auto	auto	10/100BaseTX
Fa0/17	--	disabled	1	auto	auto	10/100BaseTX
Fa0/18	--	disabled	1	auto	auto	10/100BaseTX

```
Fa0/19   Trunk to SW2   connected   trunk   a-full   a-100   10/100BaseTX
Fa0/20   Trunk to SW2   connected   trunk   a-full   a-100   10/100BaseTX
Fa0/21   --             disabled    1       auto     auto    10/100BaseTX
Fa0/22   --             disabled    1       auto     auto    10/100BaseTX
Fa0/23   --             disabled    1       auto     auto    10/100BaseTX
Fa0/24   --             disabled    1       auto     auto    10/100BaseTX
```

Note the interface description can be extremely helpful specially if the switches are configured in transparent mode, and/or the task asks for the configuration of allowed VLANs on the trunks.

Task 6

Configure Private VLANs based on the following policy:

Router	Interface	VLAN-Type	VLAN-ID
R1	F0/0	Primary	10
R2	F0/0	Community	20
R3	F0/1	Community	20
R4	F0/0	Community	30
R5	F0/1	Community	30
R6	F0/1	Isolated	40
BB1	F0/1	Isolated	40
BB2	F0/0	Isolated	40
BB3	F0/0	Isolated	40

Private-VLANs are typically seen in service provider networks, this feature addresses two major problems that the providers used to face:

1. **Number of Clients: If every client was in a VLAN of their own, the provider will be restricted to 4094 clients, which is the maximum number of VLANs on a given switch.**

2. **Routing between VLANs & IP addressing: Routing between VLANs will be a nightmare, and the number of wasted IP addresses that result from Subnetting will be enormous.**

Private-VLANs solves these two issues, with Private-VLANs a VLAN is sub-divided into sub-VLANs or sub-domains.

Private-VLANs consist of one primary, and one or more secondary VLANs, the secondary VLANs can be either Community VLANs or Isolated VLANs.

A Primary VLAN can have many Community VLANs, but it can ONLY have a Single Isolated VLAN.

Ports in a Private-VLAN:

There are three types of ports in Private-VLAN and they are as follows:

1. Promiscuous: A promiscuous port belongs to the primary VLAN; this port can communicate with all ports that are member of a secondary VLAN/s (Community and/or Isolated) that are associated with the primary VLAN that it belongs.

2. Isolated: An isolated port is a host port that belongs to an isolated secondary VLAN. The host ports that are member of a given Isolated VLAN can NOT Communicate with each other. These ports can ONLY communicate with the Port configured as Promiscuous port.

3. Community: A community port is a host port that belongs to a community Secondary VLAN. Community ports can communicate with ports in the same Community VLAN and with the port that is configured as promiscuous ports. These ports can't Communicate with other ports in other Community VLANs.

On Both Switches:

In order to configure private-vlans, the switches must be configured in Transparent mode as follows:

(config)#**vtp mode transparent**

The following commands configures the primary VLAN

(config)#**vlan 10**
(config-vlan)#**private-vlan primary**
(config-vlan)#**Exit**

The following two VLANs are defined as the community secondary VLANs, there could be many community VLANs:

(config)#**vlan 20**
(config-vlan)#**private-vlan community**

(config)#vlan 30
(config-vlan)#**private-vlan community**

There can ONLY be one isolated secondary VLAN:

(config)#vlan 40
(config-vlan)#**private-vlan isolated**

The following command associates the secondary VLANs to the primary:

(config)#vlan 10
(config-vlan)#**private-vlan association add 20,30,40**

To verify the configuration:

On Both Switches:

SWx#Show vlan private-vlan

Primary	Secondary	Type	Ports
10	20	community	
10	30	community	
10	40	isolated	

The output of the above show command displays the secondary VLANs that are created so far and the primary VLAN to which they are associated.

On SW1

The following command sets F0/1 interface in promiscuous mode, assigns the port to primary VLAN 10 and maps VLANs 20, 30 and 40 to this interface:

SW1(config)#**Int F0/1**
SW1(config-if)#**Switchport mode private-vlan promiscuous**
SW1(config-if)#**Switchport private-vlan mapping 10 20,30,40**

The ports that belong to a given secondary VLAN must be configured in host mode. The following command sets F0/2 interface in a host mode, associates this port to VLAN 10 (The primary VLAN) and assigns this port to VLAN 20 which was configured as a community secondary VLAN earlier:

SW1(config-if)#**Int F0/2**
SW1(config-if)#**Switchport mode private-vlan host**

SW1(config-if)#**Switchport private-vlan host-association 10 20**

The following command sets F0/4 interface in a host mode, associates this port to VLAN 10 (The primary VLAN) and assigns this port to VLAN 30 which was configured as a community secondary VLAN earlier:

SW1(config-if)#**Int F0/4**
SW1(config-if)#**Switchport mode private-vlan host**
SW1(config-if)#**switchport private-vlan host-association 10 30**

The following command sets F0/11 and F0/12 interfaces in a host mode, associates these ports to VLAN 10 (The primary VLAN) and assigns these ports to VLAN 40 which was configured as an isolated secondary VLAN earlier:

SW1(config)#**Int range F0/11-12**
SW1(config-if)#**Switchport mode private-vlan host**
SW1(config-if)#**Switchport private-vlan host-association 10 40**

To verify the configuration:

On SW1

SW1#Sh vlan pri

```
Primary Secondary Type              Ports
------- --------- ----------------- -----------
10      20        community         Fa0/1, Fa0/2
10      30        community         Fa0/1, Fa0/4
10      40        isolated          Fa0/1, Fa0/11, Fa0/12
```

On SW2

SW2(config)#**Int F0/3**
SW2(config-if)#**Switchport mode private-vlan host**
SW2(config-if)#**Switchport private-vlan host-association 10 20**

SW2(config)#**Int F0/5**
SW2(config-if)#**Switchport mode private-vlan host**
SW2(config-if)#**Switchport private-vlan host-association 10 30**

SW2(config)#**Int range F0/6 , F0/11**
SW2(config-if)#**Switchport mode private-vlan host**
SW2(config-if)#**switchport private-vlan host-association 10 40**

To verify the configuration:

On SW2

SW2#Show vlan private-vlan

Primary	Secondary	Type	Ports
10	20	community	Fa0/3
10	30	community	Fa0/5
10	40	isolated	Fa0/6, Fa0/11

To test the configuration:

On R1

R1#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Note R1 is able to ping all routers because it is configured to be in promiscuous mode, this interface can be thought of as the default gateway.

On R2

R2#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R2#Ping 200.1.1.3

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note R2 is able to ping R1 which is the port in the primary VLAN and R3 which is in the same community VLAN.

```
R2#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.9

Type escape sequence to abort.
```

Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note R2 was NOT able to ping the other routers because they are NOT in primary or in the same community secondary VLAN.

On R3

R3#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Note R3 is able to ping R1 which is the port in primary VLAN and the router in its own community secondary VLAN, which is R2.

R3#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:

.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.10, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.10, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note R3 can NOT ping the other routers because they are in another secondary VLAN.

On R4

R4#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R4#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Note R4 is able to ping R1 which is the port in primary VLAN and the router in its own community secondary VLAN, which is R5.

R4#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note R4 can NOT ping the other routers because they are in another secondary VLAN.

On R5

R5#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R5#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Note R5 is able to ping R1 which is the port in primary VLAN and the router in its own community secondary VLAN (R2).

R5#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:

.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note R5 can NOT ping the other routers because they are in another secondary VLAN.

On R6

R6#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Note R6 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though BB1, BB2 and BB3 are in the same VLAN, but remember that the VLAN is defined as isolated; the hosts in isolated VLAN do NOT have reachability to each other.

R6#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:

.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

On BB1

BB1#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Note BB1 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R6, BB2 and BB3 are in the same VLAN, but remember that the VLAN is defined as an isolated secondary VLAN; the hosts in isolated VLAN do NOT have reachability to each other.

BB1#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB1#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB1#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB1#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB1#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB1#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

On BB2

BB2#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Note BB2 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R6, BB1 and BB3 are in the same VLAN, but remember that the VLAN is defined as an isolated secondary VLAN; the hosts in isolated VLAN do NOT have reachability to each other.

BB2#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB2#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB2#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:

.....
Success rate is 0 percent (0/5)

BB2#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB2#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB2#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB2#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

On BB3

BB3#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Note BB3 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R6, BB1 and BB2 are in the same VLAN, but remember that the VLAN is defined as an isolated secondary VLAN; the hosts in isolated VLAN do NOT have reachability to each other.

BB3#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB3#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB3#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB3#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB3#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB3#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

BB3#Ping 200.1.1.8

Type escape sequence to abort.

```
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Task 7

Reconfigure the IP addressing of the hosts that belong to the two community secondary VLANs based on the following chart and provide InterVlan routing between them: The hosts in the other secondary VLANs should still be able to reach the host in the primary VLAN. You can use static routes and any IP addressing to accomplish this task.

Routers	IP address	VLAN-ID
R2	202.1.1.2 /24	20
R3	202.1.1.3 /24	20
R4	203.1.1.4 /24	30
R5	203.1.1.5 /24	30

On R2

```
R2(config)#int f0/0
R2(config-if)#ip addr 202.1.1.2 255.255.255.0

R2(config)#ip route 0.0.0.0 0.0.0.0 202.1.1.100
```

On R3

```
R3(config)#int f0/1
R3(config-if)#ip addr 202.1.1.3 255.255.255.0

R3(config)#ip route 0.0.0.0 0.0.0.0 202.1.1.100
```

On R4

```
R4(config)#int f0/0
R4(config-if)#ip addr 203.1.1.4 255.255.255.0

R4(config)#ip route 0.0.0.0 0.0.0.0 203.1.1.100
```

On R5

R5(config)#int f0/1
R5(config-if)#ip addr 203.1.1.5 255.255.255.0

R5(config)#**ip route** 0.0.0.0 0.0.0.0 **203.1.1.100**

On SW1

SW1(config)#**Ip routing**

Note two IP addresses are configured under interface VLAN 10, a primary and a secondary, the primary IP address is used by VLAN 20 and the secondary is used by the hosts in VLAN 30.
The "Private-vlan mapping" command maps the secondary VLANs to their layer 3 VLAN interface, in this case VLAN 10 which is the layer 3 interface of the primary VLAN.

SW1(config)#int vlan 10
SW1(config-if)#ip addr 202.1.1.100 255.255.255.0
SW1(config-if)#ip addr 203.1.1.100 255.255.255.0 sec
SW1(config-if)#**private-vlan mapping 20,30**

With the "Private-vlan mapping" interface configuration command, secondary VLANs can be added or removed using the "Private-vlan mapping add, or Private-vlan mapping remove" interface configuration command. After this command is entered, you should get the following messages:

%PV-6-PV_MSG: Created a private vlan mapping, Primary 10, Secondary 20
%PV-6-PV_MSG: Created a private vlan mapping, Primary 10, Secondary 30

To verify the configuration:

On SW1

SW1#Show interfaces private-vlan mapping

```
Interface Secondary VLAN  Type
--------- --------------- -----------------

vlan10    20              community
vlan10    30              community
```

To test the configuration:

On R2

R2#Ping 203.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 203.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (4/5), round-trip min/avg/max = 1/1/4 ms

On BB1

BB1#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Task 8

Erase the startup config and reload the routers before proceeding to the next task.

Advanced CCIE Routing & Switching 2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

Frame-relay

Lab 1 – Hub-n-Spoke using Frame-relay map statements

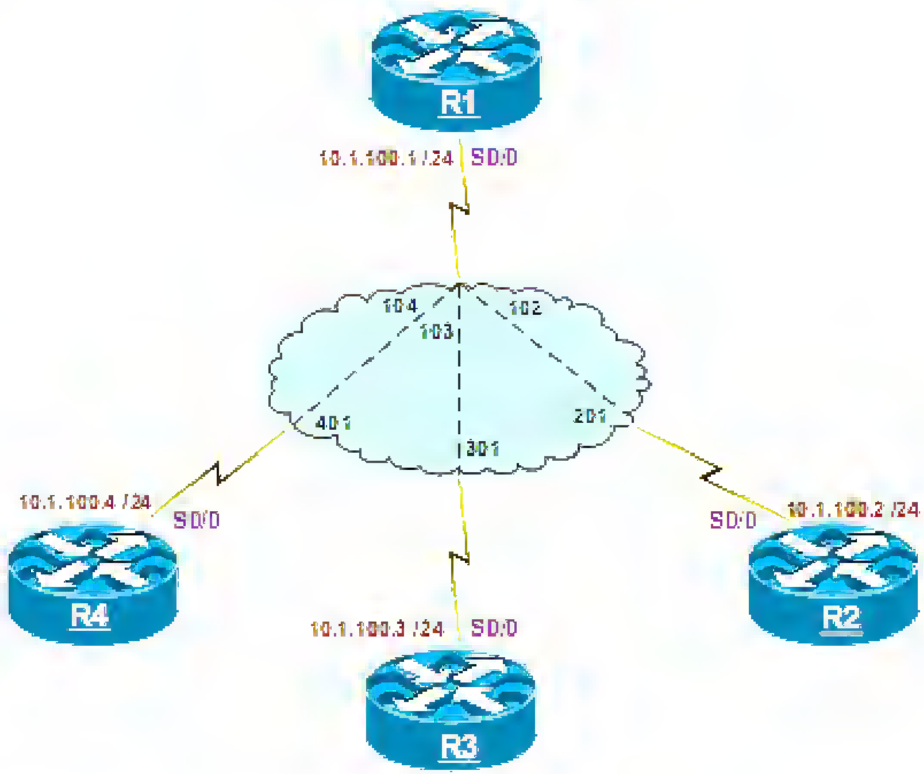

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Frame-relay interface S0/0	10.1.100.1 /24	102 103 104	R2 R3 R4
R2's Frame-relay interface S0/0	10.1.100.2 /24	201	R1
R3's Frame-relay interface S0/0	10.1.100.3 /24	301	R1
R4's Frame-relay interface S0/0	10.1.100.4 /24	401	R1

Task 1

Configure a frame-relay Hub and spoke using frame-relay map statements. Use the IP addressing in the above chart.

Disable inverse-arp such that the routers do not generate inverse-arp request packets, and ensure that only the assigned DLCIs are used and mapped, these mappings should be as follows:

- On R1: 102, 103 and 104 should be mapped to R2, R3 and R4 respectively.
- On R2, R3 and R4: DLCIs 201, 301 and 401 should be used on R2, R3 and R4 respectively for their mapping to R1 (The hub).

In the future Eigrp routing protocol will be running on these routers, ensure that the routers can handle the Multicast traffic generated by the Eigrp routing protocol. DO NOT configure any sub-interface/s to accomplish this task.

On R1

R1(config)#Int S0/0
R1(config-if)#Ip address 10.1.100.1 255.255.255.0
R1(config-if)# Encapsulation frame
R1(config-if)# Frame-relay map ip 10.1.100.2 102 broadcast
R1(config-if)# Frame-relay map ip 10.1.100.3 103 broadcast
R1(config-if)#Frame-relay map ip 10.1.100.4 104 broadcast
R1(config-if)#NO frame-relay inverse-arp
R1(config-if)#NO shut

To verify the configuration:

On R1

R1#Show frame map

Serial0/0 (up): ip 10.1.100.2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, inactive
Serial0/0 (up): ip 10.1.100.3 dlci 103(0x67,0x1870), static,
broadcast,
CISCO, status defined, inactive
Serial0/0 (up): ip 10.1.100.4 dlci 104(0x68,0x1880), static,
broadcast,
CISCO, status defined, inactive

Note you may see DLCIs 105 and 106 mapped to 0.0.0.0 IP address, these dynamic mappings may not affect Unicast traffic, but they will definitely affect Multicast and/or Broadcast traffic, therefore, they should be removed from the mapping table. The "clear frame-relay inarp" command will NOT have any effect on these entries, whereas, saving the configuration and then reloading the routers will definitely clear the 0.0.0.0 mappings. Another way to clear the "0.0.0.0" mapping is to remove the encapsulation and reconfigure the encapsulation back again, but once the encapsulation is removed, the frame-relay maps are also removed, therefore, the frame-relay maps must be re-entered.

On R1

R1#Wr

R1#Reload

R1#Show frame map

Serial0/0 (up): ip 10.1.100.2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, inactive

Serial0/0 (up): ip 10.1.100.3 dlci 103(0x67,0x1870), static,
broadcast,
CISCO, status defined, inactive

Serial0/0 (up): ip 10.1.100.4 dlci 104(0x68,0x1880), static,
broadcast,
CISCO, status defined, inactive

Note the inactive status means that the problem is on the other side of the VC, in this case the other end of these VCs are not configured yet, and once they are configured, the status should transition to active state.

The following explains the output of the "Show frame-relay map" command:
In this case the first mapping is analyzed:

Serial0/0 (up): ip 10.1.100.2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, inactive

Serial0/0 (up): ip 10.1.100.2:

This is the interface through which IP 10.1.100.2 is found.

Dlci 102(0x66,0x1860), static:

Dlci 102, this is the local DLCI that is mapped to 10.1.100.2. In the parentheses you find 2 Hexadecimal values, in this case: 0x66, 0x1860:
If the Hexadecimal 0x66 is converted to decimal, the result is 102, which is the local DLCI number.
The second Hexadecimal value 0x1860, indicates how the DLCI is split into two sections within the Frame-relay header, remember that the first 6 bits (The most significant 6 bits) are in the first byte and the last 4 bits of the DLCI, is found in the beginning of the second byte of the Frame-relay frame, as follows:

Convert 0x1860 to Binary:

1	8	6	0
0001	1 0 0 0	0110	0000

Take the most significant 6 bits, in this case: 000110

Take the most significant 4 bits of the second byte, in this case: 011000

Note the most significant 6 bits of the first byte and the most significant 4 bits of the second byte are concatenated into a 10 bit value, as follows:

00011001 1000
If the above binary number is converted to decimal, you should see 102.

On R2

R2(config)#Int S0/0
R2(config-if)#Ip address 10.1.100.2 255.255.255.0
R2(config-if)#Encapsulation frame
R2(config-if)#Frame-relay map ip 10.1.100.1 201 broadcast
R2(config-if)#**NO frame-relay inverse-arp**
R2(config-if)#**NO** shut

To verify the configuration:

On R2

R2#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R2#Show frame map

Serial0/0 (up): ip 10.1.100.1 dlci 201(0xC9,0x3090), static,
broadcast,
CISCO, status defined, active

On R3

R3(config)#Int S0/0
R3(config-if)#Ip address 10.1.100.3 255.255.255.0
R3(config-if)#Encapsulation frame
R3(config-if)#Frame-relay map ip 10.1.100.1 301 broadcast
R3(config-if)#**NO frame-relay inverse-arp**
R3(config-if)#**NO** shut

To verify the configuration:

On R3

R3#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R3#Show frame map

Serial0/0 (up): ip 10.1.100.1 dlci 301(0x12D,0x48D0), static,
broadcast,
CISCO, status defined, active

On R4

R4(config)#Int S0/0
R4(config)#Ip address 10.1.100.4 255.255.255.0
R4(config)#Encapsulation frame

R4(config)#Frame-relay map ip 10.1.100.1 401 broadcast
R4(config)#**NO frame-relay inverse-arp**
R4(config)#**NO** shut

To verify the configuration:

On R4

R4#Show frame map

Serial0/0 (up): ip 10.1.100.1 **dlci** 401(0x191,0x6410), **static,**
broadcast,
CISCO, status defined, active

R4#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

Task 2

Ensure that every router can ping every IP address connected to the cloud. When configuring this task, ensure that the hub router does NOT receive redundant routing traffic.

On R1

To test the existing configuration:

R1#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note in a multipoint frame-relay configuration, two conditions must be met before an IP address is reachable:

A. **The destination IP address must be in the routing table with a valid next hop.**

B. **There must be a frame-relay mapping for that destination.**

In this case the destination IP address is in the routing table, but the frame-relay mapping is missing. Configure the frame-relay mapping as follows:

On R1

R1(config)#Interface S0/0
R1(config-if)#**Frame-relay map ip** 10.1.100.1 102

Note there is no need to add the "broadcast" keyword for this configuration.

To verify the configuration:

On R1

R1#Show frame map

Serial0/0 (up): ip 10.1.100.1 dlci 102(0x66,0x1860), static,
CISCO, status defined, active

Serial0/0 (up): ip 10.1.100.2 dlci 102(0x66,0x1860), static,
broadcast,
CISCO, status defined, active

Serial0/0 (up): ip 10.1.100.3 dlci 103(0x67,0x1870), static,
broadcast,
CISCO, status defined, active

Serial0/0 (up): ip 10.1.100.4 dlci 104(0x68,0x1880), static,
broadcast,
CISCO, status defined, active

R1#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/124 ms

On R2

R2(config-if)#Interface S0/0
R2(config-if)#**Frame-relay map ip 10.1.100.3 201**
R2(config-if)#**Frame-relay map ip 10.1.100.4 201**
R2(config-if)#**Frame-relay map ip 10.1.100.2 201**

To verify the configuration:

On R2

R2#Ping 10.1.100.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

On R3

R3(config)#Interface S0/0
R3(config-if)#**Frame-relay map ip 10.1.100.2 301**
R3(config-if)#**Frame-relay map ip 10.1.100.4 301**
R3(config-if)#**Frame-relay map ip 10.1.100.3 301**

To verify the configuration:

On R3

R3#Ping 10.1.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

On R4

R4(config)#Interface S0/0
R4(config-if)#**Frame-relay map ip** 10.1.100.2 401
R4(config-if)#**Frame-relay map ip** 10.1.100.3 401
R4(config-if)#**Frame-relay map ip** 10.1.100.4 401

Note when configuring the frame-relay mapping from one spoke to another spoke, the "broadcast" keyword should not be used, if this keyword is used, the hub router will receive redundant routing traffic. This can be verified by running RIPv2 and performing a "debug ip rip" command on the hub router.

To verify the configuration:

On R4

R4#Ping 10.1.100.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/116 ms

R4#Ping 10.1.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/116 ms

R4#Ping 10.1.100.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/120 ms

On R3

R3#Ping 10.1.100.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/116 ms

Task 3

Configure the routers such that the LMI status inquiries are sent every 5 seconds and Full Status LMI requests are sent every 3 cycles instead of 6.

By default frame-relay routers generate LMI Status inquiries every 10 seconds, and a full status inquiry every 6th cycle (Every 60 seconds). The interval for status inquiries can be changed using the "Keepalive" command, whereas, the "Frame-relay lmi-n391dte" command can be used to change the interval for the complete status inquiries.

NOTE the output of the following debug command reveals the status inquiries and full status inquiries:

On R1

R1#Debug frame lmi

*Nov 24 19:59:57.407: Serial0/0(out): StEnq, myseq 125, yourseen 124, DTE up
*Nov 24 19:59:57.407: datagramstart = 0x3F401ED4, datagramsize = 14
*Nov 24 19:59:57.407: FR encap = 0x00010308
*Nov 24 19:59:57.407: 00 75 95 01 01 01 03 02 7D 7C

*Nov 24 19:59:57.411: Serial0/0(in): Status, myseq 125, pak size 14
*Nov 24 19:59:57.411: RT IE 1, length 1, **type 1**
*Nov 24 19:59:57.411: KA IE 3, length 2, yourseq 125, myseq 125

```
*Nov 24 20:00:07.407: Serial0/0(out): StEnq, myseq 126, yourseen 125, DTE up
*Nov 24 20:00:07.407: datagramstart = 0x3F6B0294, datagramsize = 14
*Nov 24 20:00:07.407: FR encap = 0x00010308
*Nov 24 20:00:07.407: 00 75 95 01 01 01 03 02 7E 7D

*Nov 24 20:00:07.411: Serial0/0(in): Status, myseq 126, pak size 14
*Nov 24 20:00:07.411: RT IE 1, length 1, type 1
*Nov 24 20:00:07.411: KA IE 3, length 2, yourseq 126, myseq 126

*Nov 24 20:00:17.407: Serial0/0(out): StEnq, myseq 127, yourseen 126, DTE up
*Nov 24 20:00:17.407: datagramstart = 0x3F400C14, datagramsize = 14
*Nov 24 20:00:17.407: FR encap = 0x00010308
*Nov 24 20:00:17.407: 00 75 95 01 01 01 03 02 7F 7E
*Nov 24 20:00:17.407:
*Nov 24 20:00:17.411: Serial0/0(in): Status, myseq 127, pak size 14
*Nov 24 20:00:17.411: RT IE 1, length 1, type 1
*Nov 24 20:00:17.411: KA IE 3, length 2, yourseq 127, myseq 127

*Nov 24 20:00:27.407: Serial0/0(out): StEnq, myseq 128, yourseen 127, DTE up
*Nov 24 20:00:27.407: datagramstart = 0x3F6AF394, datagramsize = 14
*Nov 24 20:00:27.407: FR encap = 0x00010308
*Nov 24 20:00:27.407: 00 75 95 01 01 01 03 02 80 7F
*Nov 24 20:00:27.407:
*Nov 24 20:00:27.411: Serial0/0(in): Status, myseq 128, pak size 14
*Nov 24 20:00:27.411: RT IE 1, length 1, type 1
*Nov 24 20:00:27.411: KA IE 3, length 2, yourseq 128, myseq 128

*Nov 24 20:00:37.407: Serial0/0(out): StEnq, myseq 129, yourseen 128, DTE up
*Nov 24 20:00:37.407: datagramstart = 0x3F644ED4, datagramsize = 14

*Nov 24 20:00:37.407: FR encap = 0x00010308
*Nov 24 20:00:37.407: 00 75 95 01 01 01 03 02 81 80
*Nov 24 20:00:37.407:
*Nov 24 20:00:37.411: Serial0/0(in): Status, myseq 129, pak size 14
*Nov 24 20:00:37.411: RT IE 1, length 1, type 1
*Nov 24 20:00:37.411: KA IE 3, length 2, yourseq 129, myseq 129

*Nov 24 20:00:47.407: Serial0/0(out): StEnq, myseq 130, yourseen 129, DTE up
*Nov 24 20:00:47.407: datagramstart = 0x3F6B03D4, datagramsize = 14
*Nov 24 20:00:47.407: FR encap = 0x00010308
*Nov 24 20:00:47.407: 00 75 95 01 01 00 03 02 82 81

*Nov 24 20:00:47.419: Serial0/0(in): Status, myseq 130, pak size 59
*Nov 24 20:00:47.419: RT IE 1, length 1, type 0
*Nov 24 20:00:47.419: KA IE 3, length 2, yourseq 130, myseq 130
```

```
*Nov 24 20:00:47.419: PVC IE 0x7 , length 0x3 , dlci 102, status 0x2
*Nov 24 20:00:47.419: PVC IE 0x7 , length 0x3 , dlci 103, status 0x2
*Nov 24 20:00:47.419: PVC IE 0x7 , length 0x3 , dlci 104, status 0x2
*Nov 24 20:00:47.419: PVC IE 0x7 , length 0x3 , dlci 105, status 0x0
*Nov 24 20:00:47.419: PVC IE 0x7 , length 0x3 , dlci 106, status 0x0
```

Note the status inquiries are sent every 10 seconds, these messages are "type 1s", whereas, the complete status inquiries are generated by the local router every 6th cycle, these message are "type 0" messages, and when the frame-relay switch receives these messages it responds with all the DLCIs that are configured for that given router.

To change these timers:

On all routers

```
(config)#Interface S0/0
(config-if)#Keepalive 5
(config-if)#Frame-relay lmi-n391dte 3
```

To test the configuration:

Rx#Debug frame LMI

```
*Nov 24 20:13:52.411: Serial0/0(out): StEnq, myseq 221, yourseen 220, DTE up
*Nov 24 20:13:52.411: datagramstart = 0x3F6AEFD4, datagramsize = 14
*Nov 24 20:13:52.411: FR encap = 0x00010308
*Nov 24 20:13:52.411: 00 75 95 01 01 01 03 02 DD DC

*Nov 24 20:13:52.415: Serial0/0(in): Status, myseq 221, pak size 14
*Nov 24 20:13:52.415: RT IE 1, length 1, type 1
*Nov 24 20:13:52.415: KA IE 3, length 2, yourseq 221, myseq 221

*Nov 24 20:13:57.411: Serial0/0(out): StEnq, myseq 222, yourseen 221, DTE up
*Nov 24 20:13:57.411: datagramstart = 0x3F400D54, datagramsize = 14
*Nov 24 20:13:57.411: FR encap = 0x00010308
*Nov 24 20:13:57.411: 00 75 95 01 01 01 03 02 DE DD

*Nov 24 20:13:57.415: Serial0/0(in): Status, myseq 222, pak size 14
*Nov 24 20:13:57.415: RT IE 1, length 1, type 1
*Nov 24 20:13:57.415: KA IE 3, length 2, yourseq 222, myseq 222

*Nov 24 20:14:02.411: Serial0/0(out): StEnq, myseq 223, yourseen 222, DTE up
```

```
*Nov 24 20:14:02.411: datagramstart = 0x3F6AF394, datagramsize = 14
*Nov 24 20:14:02.411: FR encap = 0x00010308

*Nov 24 20:14:02.411: 00 75 95 01 01 00 03 02 DF DE

*Nov 24 20:14:02.423: Serial0/0(in): Status, myseq 223, pak size 59
*Nov 24 20:14:02.423: RT IE 1, length 1, type 0
*Nov 24 20:14:02.423: KA IE 3, length 2, yourseq 223, myseq 223
*Nov 24 20:14:02.423: PVC IE 0x7 , length 0x3 , dlci 102, status 0x2

*Nov 24 20:14:02.423: PVC IE 0x7 , length 0x3 , dlci 103, status 0x2
*Nov 24 20:14:02.423: PVC IE 0x7 , length 0x3 , dlci 104, status 0x2
*Nov 24 20:14:02.423: PVC IE 0x7 , length 0x3 , dlci 105, status 0x0
*Nov 24 20:14:02.423: PVC IE 0x7 , length 0x3 , dlci 106, status 0x0
```

Note initially the router and the frame-relay switch exchange two "type 1" inquiries, and the third message that the local router generates is a "type 0" messages which tells the switch to respond with all the DLCIs.

Task 4

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 2 – Hub-n-Spoke using Frame-relay Point-to-Point configuration

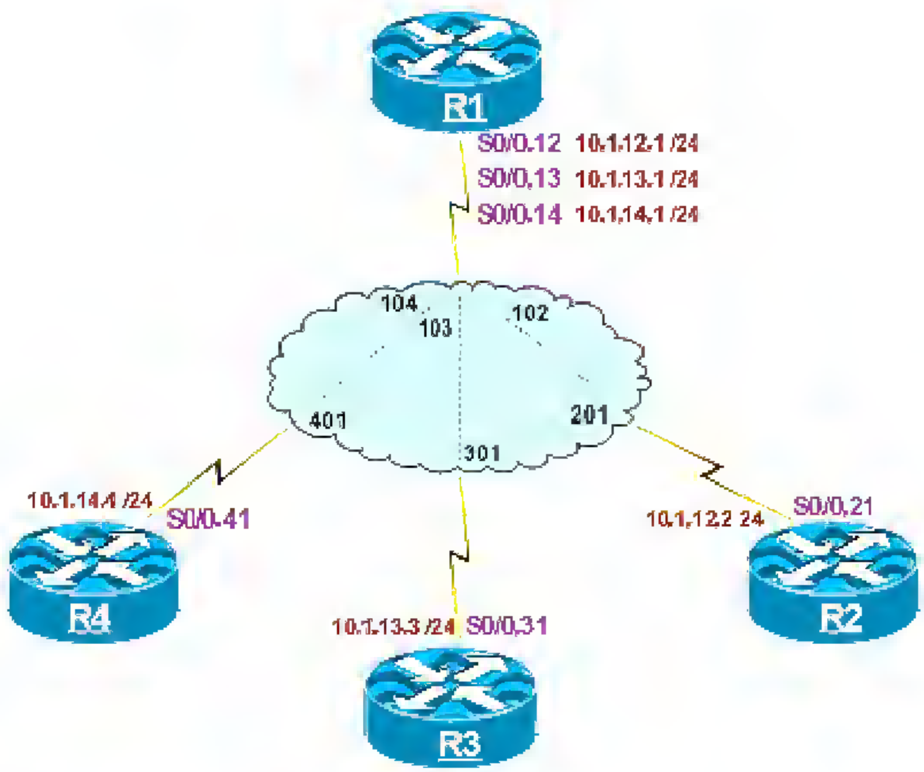

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Frame-relay interface	10.1.12.1 /24 10.1.13.1 /24 10.1.14.1 /24	102 103 104	R2 R3 R4
R2's Frame-relay interface	10.1.12.2 /24	201	R1
R3's Frame-relay interface	10.1.13.3 /24	301	R1
R4's Frame-relay interface	10.1.14.4 /24	401	R1

Task 1

Configure the routers in a hub and spoke manner using the IP addressing in the above chart.

These routers should be configured with Point-to-point sub-interface/s, and ensure that only the assigned DLCIs are used, these DLCIs should be as follows:

- **On R1:** 102, 103 and 104 should be used for connections to R2, R3 and R4 respectively.
- **On R2, R3 and R4:** DLCIs 201, 301 and 401 should be used on R2, R3 and R4 respectively for their connection to R1 (The hub).

These routers should be able to ping every IP address **within their IP address space.**

On R1

```
R1(config)#Interface S0/0
R1(config-if)# Encap frame
R1(config-if)#No shut
R1(config-if)#Exit

R1(config)#Interface S0/0.12 point-to-point
R1(config-subif)#Ip address 10.1.12.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 102
R1(config-subif)#Exit

R1(config-subif)#Interface S0/0.13 point-to-point
R1(config-subif)#Ip address 10.1.13.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 103
R1(config-subif)#Exit

R1(config-subif)#Interface S0/0.14 point-to-point
R1(config-subif)#Ip address 10.1.14.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 104
R1(config-subif)#Exit
```

To verify the configuration:

On R1

R1#Show frame map

```
Serial0/0.12 (up): point-to-point dlci, dlci 102(0x66,0x1860), broadcast
        status defined, active
Serial0/0.14 (up): point-to-point dlci, dlci 104(0x68,0x1880), broadcast
        status defined, active
Serial0/0.13 (up): point-to-point dlci, dlci 103(0x67,0x1870), broadcast
        status defined, active
```

Note when frame-relay is configured in a point-to-point manner it's important to understand the following two behaviors:

A. **There is no need to disable sending inverse-arp packets, because inverse-arp is disabled when frame-relay is configured in a point-to-point manner.**

B. **No need for frame-relay mapping/s, because there can only be another router on the other end of the PVC, therefore, all IP addresses (This includes the local router's IP address) are reachable as long as the destination IP address is in the routing table with a valid next hop IP address.**

On R2

```
R2(config)#Int S0/0
R2(config-if)#Encap frame
R2(config-if)#No shut
R2(config-if)#Exit

R2(config)#Int S0/0.21 point-to-point
R2(config-subif)#Ip address 10.1.12.2 255.255.255.0
R2(config-subif)#Frame-relay interface-dlci 201
R2(config-subif)#Exit
```

To verify and test the configuration:

On R2

```
R2#Show frame map

Serial0/0.21 (up): point-to-point dlci, dlci 201(0xC9,0x3090), broadcast
        status defined, active

R2#Ping 10.1.12.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
```

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R2#Ping 10.1.12.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

On R3

R3(config)#Int S0/0
R3(config-if)# Encap frame
R3(config-if)# No shut
R3(config-if)# Exit

R3(config)#Int S0/0.31 point-to-point
R3(config-subif)#Ip address 10.1.13.3 255.255.255.0
R3(config-subif)#Frame-relay interface-dlci 301

To verify and test the configuration:

On R3

R3#Show frame map

Serial0/0.31 (up): point-to-point dlci, dlci 301(0x12D,0x48D0), broadcast
status defined, active

R3#Ping 10.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R3#Ping 10.1.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

On R4

```
R4(config)#Int S0/0
R4(config-if)#Encap frame
R4(config-if)#No shut
R4(config-if)#Exit

R4(config)#Int S0/0.41 point-to-point
R4(config-subif)#Ip address 10.1.14.4 255.255.255.0
R4(config-subif)#Frame-relay interface-dlci 401
```

To verify and test the configuration:

On R4

R4#Show frame map

```
Serial0/0.41 (up): point-to-point dlci, dlci 401(0x191,0x6410), broadcast
         status defined, active
```

R4#Ping 10.1.14.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

R4#Ping 10.1.14.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms
```

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 3 – Mixture of Point-to-point and Multipoint frame-relay

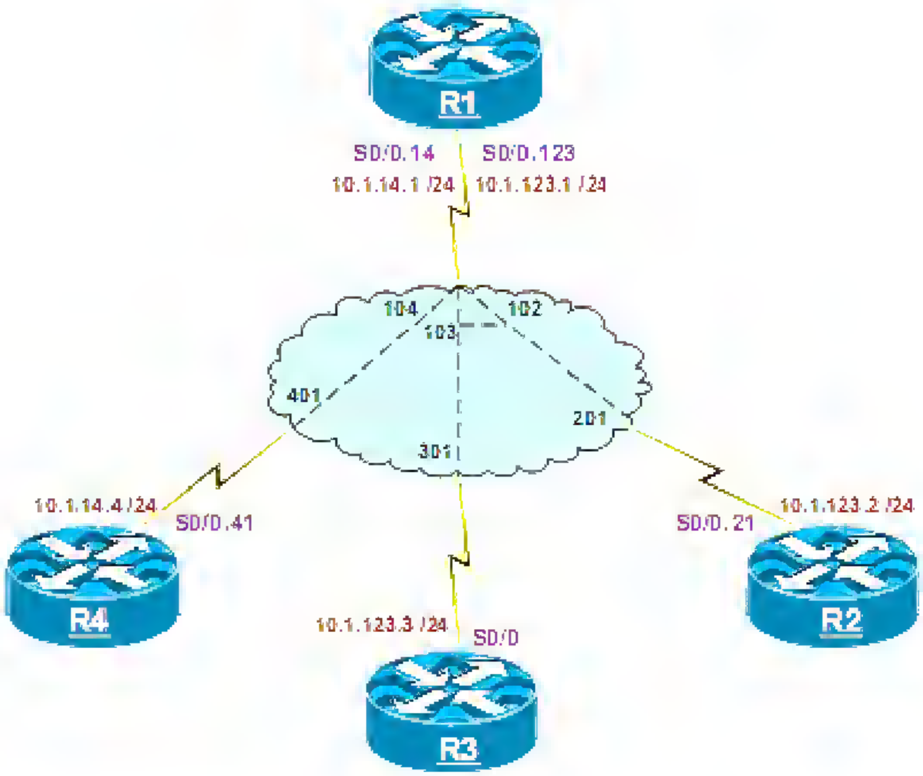

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Frame-relay interface	10.1.123.1 /24 10.1.123.1 /24 10.1.14.1 /24	102 103 104	R2 R3 R4
R2's Frame-relay interface	**10.1.123.2 /24**	**201**	**R1**
R3's Frame-relay interface	10.1.123.3 /24	301	R1
R4's Frame-relay interface	**10.1.14.4 /24**	**401**	**R1**

Task 1

Configure frame-relay on the routers as follows:

R1: This router should be configured in a point-to-point manner for it's connection to R4 and in a Multipoint manner for its connection to R2 and R3. Use the IP addressing and DLCI information in the above chart.

R2: This router should be configured in a point-to-point manner for its connection to R1. Use the IP addressing and DLCI information in the above chart.

R3: This router should be configured using its main interface for its connection to R1. Use the IP addressing and DLCI information in the above chart.

R4: This router must be configured in a point-to-point manner for its connection to R1. Use the IP addressing and DLCI information in the above chart.

Disable inverse-arp where appropriate. These routers should be able to ping every IP address within their IP address space.

On R1

R1(config)#Int S0/0
R1(config-if)# Encap frame
R1(config-if)#No shut
R1(config-if)#Exit

R1(config-subif)#Int S0/0.123 multipoint
R1(config-subif)#Ip address 10.1.123.1 255.255.255.0
R1(config-subif)#**Frame-relay map ip 10.1.123.2 102**
R1(config-subif)#**Frame-relay map ip 10.1.123.3 103**

R1(config)#Int S0/0.14 point-to-point
R1(config-subif)#Ip address 10.1.14.1 255.255.255.0
R1(config-subif)#**Frame-relay interface-dlci 104**

On R2

R2(config)#Int S0/0
R2(config-if)# Encap frame
R2(config-if)#No shut
R2(config-if)#Exit

R2(config)#Int S0/0.21 point-to-point

R2(config-subif)# Ip address 10.1.123.2 255.255.255.0
R2(config-subif)#**Frame-relay interface-dlci 201**

Note there is no need to disable sending inverse-arp, because it's disabled when a sub-interface is configured.

To test and verify the configuration:

On R2

R2#Show frame-relay map

Serial0/0.21 (up): point-to-point dlci, dlci 201(0xC9,0x3090), broadcast
status defined, active

R2#Ping 10.1.123.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.123.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

On R3

R3(config)#Int S0/0
R3(config-if)#Encap frame
R3(config-if)#Ip address 10.1.123.3 255.255.255.0
R3(config-if)#**Frame-relay map ip 10.1.123.1 301**
R3(config-if)#**Frame-relay map ip 10.1.123.2 301**
R3(config-if)#**Frame-relay map ip 10.1.123.3 301**
R3(config-if)#**No frame-relay inverse-arp**
R3(config-if)#No shut

Note inverse-arp should be Disabled because the configuration is done directly under the main interface

To verify and test the configuration:

On R3

R3#Show frame map

Serial0/0 (up): ip 10.1.123.1 dlci 301(0x12D,0x48D0), static,
CISCO, status defined, active

Serial0/0 (up): ip 10.1.123.2 dlci 301(0x12D,0x48D0), static,
CISCO, status defined, active

Serial0/0 (up): ip 10.1.123.3 dlci 301(0x12D,0x48D0), static,
CISCO, status defined, active

R3#Ping 10.1.123.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.123.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms

R3#Ping 10.1.123.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.123.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

R3#Ping 10.1.123.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.123.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

On R4

R4(config)#Int S0/0
R4(config-if)#Encap frame
R4(config-if)#No shut
R4(config-if)#Exit

R4(config)#Int S0/0.41 point-to-point
R4(config-subif)#Ip address 10.1.14.4 255.255.255.0
R4(config-subif)#Frame-relay interface-dlci 401

To verify and test the configuration:

On R4

R4#Show frame map

```
Serial0/0.41 (up): point-to-point dlci, dlci 401(0x191,0x6410), broadcast
          status defined, active

R4#Ping 10.1.14.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R4#Ping 10.1.14.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms
```

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 4 – Multipoint Frame-relay With Out Frame-relay mapping

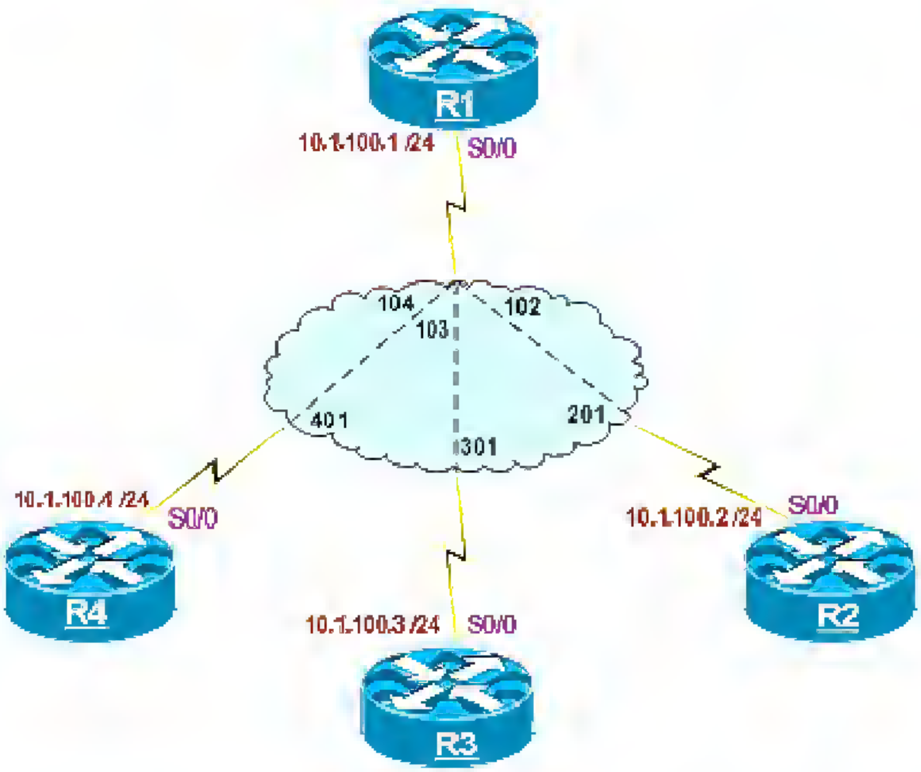

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Frame-relay interface	10.1.100.1 /24	102 103 104	R2 R3 R4
R2's Frame-relay interface	10.1.100.2 /24	201	R1
R3's Frame-relay interface	10.1.100.3 /24	301	R1
R4's Frame-relay interface	10.1.100.4 /24	401	R1

Task 1

Configure the routers in a hub and spoke manner, with R1 as the hub and R2, R3 and R4 as the spokes.
Ensure that these routers have full reachability to each other with out using the "frame-relay map" command.
Do not use PBR to accomplish this task.

In the following solution PPP is configured on the DLCIs, when PPP is configured a host route is injected into the routing table, this host route provides NLRI to the next hop IP address.

On R1

R1(config)#Interface S0/0
R1(config-if)#Encap frame-relay
R1(config-if)#**Frame-relay interface-dlci 102 ppp Virtual-Template1**
R1(config-if)#**Frame-relay interface-dlci 103 ppp Virtual-Template1**
R1(config-if)#**Frame-relay interface-dlci 104 ppp Virtual-Template1**

R1(config)#Interface Virtual-template 1
R1(config-if)#Ip address 10.1.100.1 255.255.255.0

On R2

R2(config)#Interface S0/0
R2(config-if)#Encap frame-relay
R2(config-if)#**Frame-relay interface-dlci 201 ppp Virtual-template 2**

R2(config)#Interface Virtual-template 2
R2(config-if)#Ip address 10.1.100.2 255.255.255.0

On R3

R3(config)#Interface S0/0
R3(config-if)#Encap frame-relay
R3(config-if)#**Frame-relay interface-dlci 301 ppp Virtual-template 3**

R3(config)#Interface Virtual-template 3
R3(config-if)#Ip address 10.1.100.3 255.255.255.0

On R4

R4(config)#Interface S0/0

R4(config-if)# Encap frame-relay
R4(config-if)#**Frame-relay interface-dlci 401 ppp Virtual-template 4**

R4(config)#Interface Virtual-template 4
R4(config-if)#Ip address 10.1.100.4 255.255.255.0

To verify and test the configuration:

On R1

R1#Show ip route | Inc /32

The injected host routes

```
C       10.1.100.4/32 is directly connected, Virtual-Access4
C       10.1.100.3/32 is directly connected, Virtual-Access3
C       10.1.100.2/32 is directly connected, Virtual-Access2
```

R1#Ping 10.1.100.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

R1#Ping 10.1.100.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

R1#Ping 10.1.100.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R2

R2#Show ip route | Inc /32

```
C       10.1.100.1/32 is directly connected, Virtual-Access2
```

R2#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

R2#Ping 10.1.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/116 ms

R2#Ping 10.1.100.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms

On R3

R3#Show ip route | Inc /32

C 10.1.100.1/32 is directly connected, Virtual-Access2

R3#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

R3#Ping 10.1.100.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/116 ms

R3#Ping 10.1.100.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.4, timeout is 2 seconds:
!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/116 ms

On R4

R4#Show ip route

C 10.1.100.1/32 is directly connected, Virtual-Access2

R4#Ping 10.1.100.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

R4#Ping 10.1.100.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms

R4#Ping 10.1.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 5 – Frame-relay and Authentication

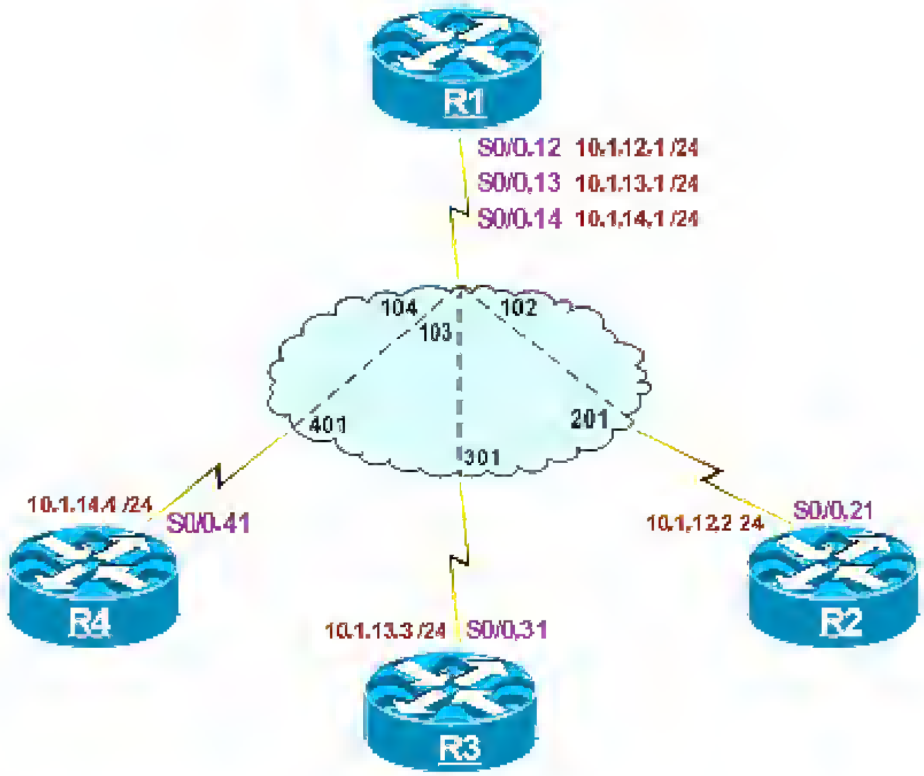

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Frame-relay interface	10.1.12.1 /24 10.1.13.1 /24 10.1.14.1 /24	102 103 104	R2 R3 R4
R2's Frame-relay interface	10.1.12.2 /24	201	R1
R3's Frame-relay interface	10.1.13.3 /24	301	R1
R4's Frame-relay interface	10.1.14.4 /24	401	R1

Task 1

Configure the routers in a hub and spoke manner using the IP addressing in the above chart.

These routers should be configured in a Point-to-Point manner as follows:

- **On R1:** DLCIs 102, 103 and 104 should be used for it's connection to R2, R3 and R4 respectively.
- **On R2, R3 and R4:** DLCIs 201, 301 and 401 should be used on R2, R3 and R4 respectively for their point-to-point frame-relay connection to R1 (The hub).

On R1

```
R1(config)#Interface S0/0
R1(config-if)#Encap frame
R1(config-if)#No shut

R1(config)#Interface S0/0.12 point-to-point
R1(config-subif)#Ip address 10.1.12.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 102
R1(config-subif)#Exit

R1(config)#Interface S0/0.13 point-to-point
R1(config-subif)# Ip address 10.1.13.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 103

R1(config)#Interface S0/0.14 point-to-point
R1(config-subif)#Ip address 10.1.14.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 104
```

To verify the configuration:

On R1

R1#Show frame map

```
Serial0/0.12 (up): point-to-point dlci, dlci 102(0x66,0x1860), broadcast
          status defined, active
Serial0/0.13 (up): point-to-point dlci, dlci 103(0x67,0x1870), broadcast
          status defined, active
Serial0/0.14 (up): point-to-point dlci, dlci 104(0x68,0x1880), broadcast
          status defined, active
```

On R2

R2(config)#Int S0/0
R2(config-if)#Encap frame
R2(config-if)#No shut

R2(config)#Int S0/0.21 point-to-point
R2(config-subif)#Ip address 10.1.12.2 255.255.255.0
R2(config-subif)#Frame-relay interface-dlci 201

To verify and test the configuration:

On R2

R2#Show frame map

Serial0/0.21 (up): point-to-point dlci, dlci 201(0xC9,0x3090), broadcast
status defined, active

R2#Ping 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

On R3

R3(config)#Int S0/0
R3(config-if)#Encap frame
R3(config-if)# No shut

R3(config)#Int S0/0.31 point-to-point
R3(config-subif)#Ip address 10.1.13.3 255.255.255.0
R3(config-subif)#Frame-relay interface-dlci 301

To verify and test the configuration:

On R3

R3#Show frame map

Serial0/0/0.31 (up): point-to-point dlci, dlci 301(0x12D,0x48D0), broadcast

```
        status defined, active

R3#Ping 10.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

On R4

```
R4(config)#Int S0/0
R4(config-if)#Encap frame
R4(config-if)#No shut

R4(config)#Int S0/0.41 point-to-point
R4(config-subif)#Ip address 10.1.14.4 255.255.255.0
R4(config-subif)#Frame-relay interface-dlci 401
```

To verify and test the configuration:

On R4

```
R4#Show frame map

Serial0/0/0.41 (up): point-to-point dlci, dlci 401(0x191,0x6410), broadcast
        status defined, active

R4#Ping 10.1.14.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

Task 2

Configure authentication on the routers as follows:

A. For R1 and R2's connection:
   R1 should send a challenge when it is called by R2.
   R2 should NOT authenticate when it is called.

The password for this authentication should be "cisco12".
This authentication should be successful even if the host name of the router is changed.

B. <u>For R1 and R3's connection:</u>
R1 should NOT authenticate when it is called.
R3 should use PAP authentication when it is called by R1.
The password for this authentication should be "cisco13".
The host name of the router should be used for this authentication.

C. <u>For R1 and R4's connection:</u>
R1 should send a challenge when it is called by R4.
R4 should use PAP authentication when it's called by R1.
The password for CHAP authentication should be "cisco", whereas, the password for PAP should be set to "ciscoPAP" and the hostname should be configured to be "R1-PAP".

<u>For R1 and R2's connection:</u>

<u>On R1</u>

R1(config)#Username R2 password 0 cisco12

R1(config)#Int S0/0.12
R1(config-if)# No IP addr
R1(config-if)#**Frame-relay interface-dlci 102 ppp virtual-template 12**

R1(config)#Int S0/0.13
R1(config-subif)#No IP address
R1(config-subif)#**Frame-relay interface-dlci 103 ppp virtual-template 13**

R1(config)#Int S0/0.14
R1(config-subif)#No IP address
R1(config-subif)#**Frame-relay interface-dlci 104 ppp virtual-template 14**

R1(config)#Int Virtual-Template12
R1(config-if)#Ip address 10.1.12.1 255.255.255.0
R1(config-if)#**ppp authentication chap callin**
R1(config-if)#**ppp chap hostname R1**

<u>On R2</u>

R2(config)#Username R1 password 0 cisco12

R2(config)#Int S0/0.21
R2(config-subif)#No IP addr
R2(config-subif)#**Frame-relay interface-dlci 201 ppp virtual-template 21**

R2(config)#Int Virtual-Template21
R2(config-if)#Ip address 10.1.12.2 255.255.255.0
R2(config-if)#**ppp chap hostname R2**

To test and verify the configuration:

On R2

R2#Debug ppp authentication

R2(config)#Int S0/0
R2(config-if)#Shut
R2(config-if)#No shut

R2(config-if)#do Ping 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:

*Nov 25 23:20:57.783: %LINK-3-UPDOWN: Interface Virtual-Access2, changed state to up

*Nov 25 23:20:59.639: Vi2 CHAP: I CHALLENGE id 17 len 23 from "R1"
*Nov 25 23:20:59.639: Vi2 CHAP: Using hostname from interface CHAP
*Nov 25 23:20:59.643: Vi2 CHAP: Using password from AAA
*Nov 25 23:20:59.643: Vi2 CHAP: O RESPONSE id 17 len 23 from "R2"
*Nov 25 23:20:59.659: Vi2 CHAP: I SUCCESS id 17 len 4.

*Nov 25 23:21:00.659: %LINEPROTO-5-UPDOWN: Line protocol on Interface Virtual-Access2, changed state to up..
Success rate is 20 percent (1/5), round-trip min/avg/max = 60/60/60 ms

The output of the above debug command shows the "Challenge" packet coming Inbound, "Response" packet going Outbound, and the "Success" coming Inbound.

For R1 and R3's connection:

On R1

R1(config)#Int Virtual-Template13
R1(config-if)#Ip address 10.1.13.1 255.255.255.0

R1(config-if)#**ppp pap sent-username R1 password 0** cisco13

On R3

R3(config)#Username R1 password 0 cisco13

R3(config)#Int S0/0.31
R3(config-subif)#No IP address
R3(config-subif)#**Frame-relay interface-dlci 301 ppp virtual-template 31**

R3(config)#Int Virtual-Template31
R3(config-if)#Ip address 10.1.13.3 255.255.255.0
R3(config-if)#**ppp authentication pap callin**

To test and verify the configuration:

On R3

R3#Debug ppp authentication

R3(config)#Int S0/0
R3(config-if)#Shut
R3(config-if)#No shut

R3(config-if)#Do Ping 10.1.13.1

*Nov 25 23:36:41.419: Vi2 PPP: Authorization required
*Nov 25 23:36:41.439: Vi2 PAP: I AUTH-REQ id 3 len 15 from "R1"
*Nov 25 23:36:41.439: Vi2 PAP: Authenticating peer R1
*Nov 25 23:36:41.439: Vi2 PPP: Sent PAP LOGIN Request
*Nov 25 23:36:41.439: Vi2 PPP: Received LOGIN Response PASS
*Nov 25 23:36:41.443: Vi2 PPP: Sent LCP AUTHOR Request
*Nov 25 23:36:41.443: Vi2 PPP: Sent IPCP AUTHOR Request
*Nov 25 23:36:41.443: Vi2 LCP: Received AAA AUTHOR Response PASS
*Nov 25 23:36:41.443: Vi2 IPCP: Received AAA AUTHOR Response PASS
*Nov 25 23:36:41.443: Vi2 PAP: O AUTH-ACK id 3 len 5
*Nov 25 23:36:41.455: Vi2 PPP: Sent IPCP AUTHOR Request

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

For R1 and R4's connection:

On R1

R1(config)#Username R4 password cisco

R1(config)#Int Virtual-Template14
R1(config)#IP address 10.1.14.1 255.255.255.0
R1(config)# **ppp authentication chap callin**
R1(config)# **ppp pap sent-username R1-PAP password 0 ciscoPAP**

On R4

R4(config)#**Username R1-PAP password ciscoPAP**
R4(config)#Username R1 password cisco

R4(config)#Int S0/0.41
R4(config-subif)#No ip address
R4(config-subif)#**Frame-relay interface-dlci 401 ppp virtual-template 41**

R4(config)#Int Virtual-Template41
R4(config-if)#IP address 10.1.14.4 255.255.255.0
R4(config-if)#**ppp authentication pap callin**

To test and verify the configuration:

On R4

R4#Debug ppp authentication

R4#(config)#Int S0/0
R4#(config-if)#Shut
R4#(config-if)#No shut

R4#(config-if)#Do Ping 10.1.14.1

```
*Mar  2 06:01:36.303: Vi1 PAP: I AUTH-REQ id 6 len 20 from "R1-PAP"
*Mar  2 06:01:36.303: Vi1 PAP: Authenticating peer R1-PAP
*Mar  2 06:01:36.307: Vi1 PPP: Sent PAP LOGIN Request
*Mar  2 06:01:36.311: Vi1 PPP: Received LOGIN Response PASS
*Mar  2 06:01:36.311: Vi1 CHAP: I CHALLENGE id 6 len 23 from "R1"
*Mar  2 06:01:36.315: Vi1 CHAP: O RESPONSE id 6 len 23 from "R4"
*Mar  2 06:01:36.319: Vi1 LCP: Received AAA AUTHOR Response PASS.!
```

```
Success rate is 20 percent (1/5), round-trip min/avg/max = 56/56/56 ms

*Mar  2 06:01:36.319: Vi1 IPCP: Received AAA AUTHOR Response PASS
*Mar  2 06:01:36.319: Vi1 PAP: O AUTH-ACK id 6 len 5
*Mar  2 06:01:36.339: Vi1 CHAP: I SUCCESS id 6 len 4
*Mar  2 06:01:36.343: Vi1 PPP: Sent IPCP AUTHOR Request
```

Task 3

Erase the startup config and reload the routers before proceeding to the next lab

Lab 6 – Frame-relay End-to-End Keepalive

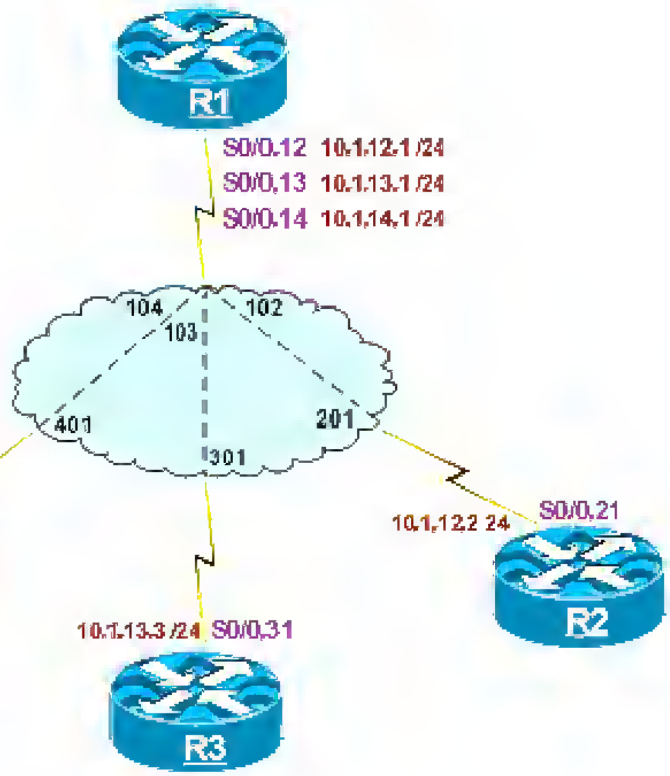

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Frame-relay interface	10.1.12.1 /24 10.1.13.1 /24 10.1.14.1 /24	102 103 104	R2 R3 R4
R2's Frame-relay interface	10.1.12.2 /24	201	R1
R3's Frame-relay interface	10.1.13.3 /24	301	R1
R4's Frame-relay interface	10.1.14.4 /24	401	R1

Task 1

Configure the routers in a hub and spoke manner using the IP addressing in the above chart.

These routers should be configured in a Point-to-Point manner as follows:

- **On R1:** DLCIs 102, 103 and 104 should be used for it's connection to R2, R3 and R4 respectively.

- **On R2, R3 and R4:** DLCIs 201, 301 and 401 should be used on R2, R3 and R4 respectively for their point-to-point frame-relay connection to R1 (The hub).

On R1

```
R1(config)#Interface S0/0
R1(config-if)# Encap frame
R1(config-if)#No shut

R1(config)#Interface S0/0.12 point-to-point
R1(config-subif)# Ip address 10.1.12.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 102

R1(config)#Interface S0/0.13 point-to-point
R1(config-subif)#Ip address 10.1.13.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 103

R1(config)#Interface S0/0.14 point-to-point
R1(config-subif)#Ip address 10.1.14.1 255.255.255.0
R1(config-subif)#Frame-relay interface-dlci 104
```

To verify the configuration:

On R1

R1#Show frame map

```
Serial0/0.12 (up): point-to-point dlci, dlci 102(0x66,0x1860), broadcast
          status defined, active
Serial0/0.13 (up): point-to-point dlci, dlci 103(0x67,0x1870), broadcast
          status defined, active
Serial0/0.14 (up): point-to-point dlci, dlci 104(0x68,0x1880), broadcast
          status defined, active
```

On R2

R2config-subif)# Int S0/0
R2config-if)#Encap frame
R2config-if)#No shut

R2config)#Int S0/0.21 point-to-point
R2config-subif)#Ip address 10.1.12.2 255.255.255.0
R2config-subif)#Frame-relay interface-dlci 201

To verify and test the configuration:

On R2

R2#Show frame map

Serial0/0.21 (up): point-to-point dlci, dlci 201(0xC9,0x3090), broadcast
status defined, active

R2#Ping 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

On R3

R3config-subif)#Int S0/0
R3config-if)#Encap frame
R3config-if)#No shut

R3config-subif)#Int S0/0.31 point-to-point
R3config-subif)#Ip address 10.1.13.3 255.255.255.0
R3config-subif)#Frame-relay interface-dlci 301

To verify and test the configuration:

On R3

R3#Show frame map

Serial0/0.31 (up): point-to-point dlci, dlci 301(0x12D,0x48D0), broadcast

```
        status defined, active

R3#Ping 10.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

On R4

```
R4config-subif)#Int S0/0
R4config-if)#Encap frame
R4config-if)#No shut

R4config-subif)#Int S0/0.41 point-to-point
R4config-subif)#Ip address 10.1.14.4 255.255.255.0
R4config-subif)#Frame-relay interface-dlci 401
```

To verify and test the configuration:

On R4

```
R4#Show frame map

Serial0/0.41 (up): point-to-point dlci, dlci 401(0x191,0x6410), broadcast
        status defined, active

R4#Ping 10.1.14.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

Task 2

Configure Frame-relay end-to-end keepalives on R1 and R2, these routers should be configured in bidirectional mode using the default values.

Routers depend on the LMIs to maintain the status of an active connection, since

The intermediate switches in the cloud may not support NNI LMIs, FREEK can be used to provide the local router with the status of the remote end. FREEK accomplishes this by providing an end to end keepalive, this keepalive runs on the data DLCI (16 – 997) and not the LMI DLCI (Cisco LMI uses DLCI 1023, and Q933a and ANSI uses DLCI 0).

FREEK maintains two internal keepalives:

- The first one is used to send out keepalive requests and to handle responses to the requests; this is considered the send side.
- The second one is to handle and reply to the requests; this is considered the receive side.

At the send side when the timer expires, the send side transmits a keepalive and waits for a reply. When the send side receives the reply before the timer expires a frame-relay keepalive is recorded. If the timer expires and no keepalives are received, an error event is recorded.
If a sufficient number of error events are observed, the PVC will transition to a down state. The number of events necessary to change the status from up to down is known as event window.
Some of the parameters and values can be changed as follows:

Frame-relay end-to-end keepalive [send | receive] error-threshold

This command configures the number of frame-relay end-to-end keepalive errors that must occur in the event window before the interface goes down. Default is 2, and the maximum number is 32.

Frame-relay end-to-end keepalive [send | receive] success-events
This command configures the number of frame-relay end-to-end keepalive successes that must occur before the interface comes up. Default is 2, and the maximum number is 32.

Frame-relay end-to-end keepalive [send | receive] timer
This command configures end to end keepalive timers; this can be configured for send or receive side

Frame-relay end-to-end keepalive event-window
This command tells the IOS to keep track of x number of most recent events.

On R1

R1(config)#Map-class frame-relay TST12
R1(config-map-class)#frame-relay end-to-end keepalive mode bidirectional

R1(config)#Int Serial0/0.12 point-to-point
R1(config-subif)#frame-relay interface-dlci 102
R1(config-fr-dlci)#class TST12

To verify the configuration:

On R1

R1#Show frame-relay end-to-end keepalive interface S0/0.12

End-to-end Keepalive Statistics for Interface Serial0/0.21 (Frame Relay DTE)

DLCI = 201, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 3, Receive Sequence Number: 4
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 6, Total Observed Errors: 0
Monitored Events: 3, Monitored Errors: 0
Successive Successes: 3, End-to-end VC Status: UP

RECEIVE SIDE STATISTICS

Send Sequence Number: 3, Receive Sequence Number: 2
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 5, Total Observed Errors: 0
Monitored Events: 3, Monitored Errors: 0
Successive Successes: 3, End-to-end VC Status: UP

On R2

R2(config)#map-class frame-relay TEST
R2(config-map-class)#frame-relay end-to-end keepalive mode bidirectional

R2(config)#interface Serial0/0.21 point-to-point
R2(config-subif)#frame interface-dlci 201
R2(config-fr-dlci)#class TEST

To verify the configuration:

On R2

R2#Show frame-relay end-to-end keepalive interface S0/0.21

End-to-end Keepalive Statistics for Interface Serial0/0.21 (Frame Relay DTE)

DLCI = 201, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 4, Receive Sequence Number: 3
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 6, Total Observed Errors: 0
Monitored Events: 3, Monitored Errors: 0
Successive Successes: 3, End-to-end VC Status: UP

RECEIVE SIDE STATISTICS

Send Sequence Number: 3, Receive Sequence Number: 2
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 5, Total Observed Errors: 0
Monitored Events: 3, Monitored Errors: 0
Successive Successes: 3, End-to-end VC Status: UP

To test the configuration:

On R2

R2(config)#Int S0/0.21
R2(config-subif)#Shut

On R1

R1#Show frame end keep inter S0/0.12

End-to-end Keepalive Statistics for Interface Serial0/0.12 (Frame Relay DTE)

DLCI = 102, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK DOWN)

SEND SIDE STATISTICS

Send Sequence Number: 29, Receive Sequence Number: 28
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 42, Total Observed Errors: 12
Monitored Events: 3, Monitored Errors: 1
Successive Successes: 0, End-to-end VC Status: UP

RECEIVE SIDE STATISTICS

Send Sequence Number: 28, Receive Sequence Number: 27
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 39, Total Observed Errors: 9
Monitored Events: 3, Monitored Errors: 2
Successive Successes: 0, End-to-end VC Status: DOWN

R1#Show ip int brie

Interface	IP-Address	OK? Method	Status	Protocol
FastEthernet0/0	unassigned	YES unset	administratively down	down
FastEthernet0/1	unassigned	YES unset	administratively down	down
Serial0/0	**unassigned**	**YES unset**	**up**	**up**
Serial0/0.12	**10.1.12.1**	**YES manual**	**down**	**down**
Serial0/0.13	10.1.13.1	YES manual	up	up
Serial0/0.14	10.1.14.1	YES manual	up	up

Note the default configured error threshold is 2, therefore, when R1 did not receive two replies within three events, it's sub-interface S0/0.12 transitioned into down/down state. But the main interface (S0/0), is still in up/up state.

To test the success events:

On R2

R2(config)#Interface S0/0.21
R2(config-subif)#No shut

On R1

R1#Sh frame end keep inter S0/0.12

End-to-end Keepalive Statistics for Interface Serial0/0.12 (Frame Relay DTE)

DLCI = 102, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 105, Receive Sequence Number: 30
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 119, Total Observed Errors: 87
Monitored Events: 0, Monitored Errors: 0
Successive Successes: 0, End-to-end VC Status: UP

RECEIVE SIDE STATISTICS

Send Sequence Number: 30, Receive Sequence Number: 29
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 90, Total Observed Errors: 58
Monitored Events: 0, Monitored Errors: 0
Successive Successes: 0, End-to-end VC Status: UP

Note after three success events in a row, the sub-interface is transitioned into up up state.

Task 3

Configure Frame-relay end-to-end keepalives for the VC that connects R1 to R3. R1 should be configured in request mode whereas R3 should be configured in reply mode using the default values.

On R1

R1(config)#Map-class frame-relay TST13
R1(config-map-class)#**frame-relay end-to-end keepalive mode request**

R1(config)#interface Serial0/0.13 point-to-point
R1(config-subif)#frame-relay interface-dlci 103
R1(config-fr-dlci)#class TST13

To verify the configuration:

On R1

R1#Show frame-relay end-to-end keepalive interface S0/0.13

End-to-end Keepalive Statistics for Interface Serial0/0.13 (Frame Relay DTE)

DLCI = 103, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 255, Receive Sequence Number: 1
Configured Event Window: 3, Configured Error Threshold: 2

```
Total Observed Events: 7,          Total Observed Errors: 4
Monitored Events: 2,               Monitored Errors: 0
Successive Successes: 2,           End-to-end VC Status: UP
```

On R3

```
R3(config)#map-class frame-relay TST31
R3(config-map-class)#frame-relay end-to-end keepalive mode reply

R3(config)#interface Serial0/0.31 point-to-point
R3(config-subif)#frame-relay interface-dlci 301
R3(config-fr-dlci)#class TST31
```

To verify the configuration:

On R3

R3#Show frame end-to-end keepalive Interface S0/0.31

```
End-to-end Keepalive Statistics for Interface Serial0/0.31 (Frame Relay DTE)

DLCI = 301, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

RECEIVE SIDE STATISTICS

Send Sequence Number: 15,          Receive Sequence Number: 14
Configured Event Window: 3,        Configured Error Threshold: 2
Total Observed Events: 17,         Total Observed Errors: 0
Monitored Events: 3,               Monitored Errors: 0
Successive Successes: 3,           End-to-end VC Status: UP
```

To test the configuration:

On R1

```
R1(config)#Int S0/0.13
R1(config-subif)#Shut
```

On R3

R3#Show frame end-to-end keepalive Interface S0/0.31

```
End-to-end Keepalive Statistics for Interface Serial0/0.31 (Frame Relay DTE)
```

DLCI = 301, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK DOWN)

RECEIVE SIDE STATISTICS

Send Sequence Number: 24, Receive Sequence Number: 23
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 31, Total Observed Errors: 5
Monitored Events: 3, Monitored Errors: 2
Successive Successes: 0, End-to-end VC Status: DOWN

To test the success events:

On R1

R1(config)#Int S0/0.13
R1(config-subif)No shut

On R3

R3#Show frame end-to-end keepalive Interface S0/0.31

End-to-end Keepalive Statistics for Interface Serial0/0.31 (Frame Relay DTE)

DLCI = 301, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

RECEIVE SIDE STATISTICS

Send Sequence Number: 26, Receive Sequence Number: 25
Configured Event Window: 3, Configured Error Threshold: 2
Total Observed Events: 42, Total Observed Errors: 14
Monitored Events: 0, Monitored Errors: 0
Successive Successes: 0, End-to-end VC Status: UP

Note the sub-interface S0/0.31 on R3 transitioned into up/up state.

Task 4

Configure Frame-relay end-to-end keepalives for the VC that connects R1 to R4. These two routers should be configured in **bidirectional** mode using the following policy:

If these routers have three errors within 5 events, the sub-interface should transition into down/down state, and if they have four success events in a row, the sub-interface should transition into up/up state. Ensure that the keepalives are exchanged every 20 seconds.

On R1

R1(config)#Map-class frame-relay TST14
R1(config-map-class)#**frame-relay end-to-end keepalive mode bidirectional**
R1(config-map-class)#**frame-relay end-to-end keepalive event-window recv 5**
R1(config-map-class)# **frame-relay end-to-end keepalive event-window send 5**

R1(config-map-class)#**frame-relay end-to-end keepalive error-threshold recv 3**
R1(config-map-class)#**frame-relay end-to-end keepalive error-threshold send 3**

R1(config-map-class)#**frame-relay end-to-end keepalive success-events recv 4**
R1(config-map-class)#**frame-relay end-to-end keepalive success-events send 4**

R1(config-map-class)#**frame-relay end-to-end keepalive timer recv 20**
R1(config-map-class)#**frame-relay end-to-end keepalive timer send 20**

R1(config)#Int Serial0/0.14 point-to-point
R1(config-subif)#frame-relay interface-dlci 104
R1(config-subif)#class TST14

To verify the configuration:

On R1

R1#Show frame-relay end-to-end keepalive interface S0/0.14

End-to-end Keepalive Statistics for Interface Serial0/0.14 (Frame Relay DTE)

DLCI = 104, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 7, Receive Sequence Number: 8
Configured Event Window: 5, Configured Error Threshold: 3
Total Observed Events: 19, Total Observed Errors: 9
Monitored Events: 5, Monitored Errors: 0
Successive Successes: 5, End-to-end VC Status: UP

RECEIVE SIDE STATISTICS

Send Sequence Number: 9,   Receive Sequence Number: 8
Configured Event Window: 5,   Configured Error Threshold: 3
Total Observed Events: 19,   Total Observed Errors: 8
Monitored Events: 5,   Monitored Errors: 0
Successive Successes: 5,   End-to-end VC Status: UP

On R4

R4(config)#Map-class frame-relay TST41
R4(config-map-class)#frame-relay end-to-end keepalive mode bidirectional

R4(config-map-class)#frame-relay end-to-end keepalive event-window recv 5
R4(config-map-class)#frame-relay end-to-end keepalive event-window send 5

R4(config-map-class)#frame-relay end-to-end keepalive error-threshold recv 3
R4(config-map-class)#frame-relay end-to-end keepalive error-threshold send 3

R4(config-map-class)#frame-relay end-to-end keepalive success-events recv 4
R4(config-map-class)#frame-relay end-to-end keepalive success-events send 4

R4(config-map-class)#frame-relay end-to-end keepalive timer recv 20
R4(config-map-class)#frame-relay end-to-end keepalive timer send 20

R4(config)#Int Serial0/0.41 point-to-point
R4(config-subif)#frame-relay interface-dlci 401
R4(config-fr-dlci)#class TST41

To verify the configuration:

On R4

R4#Show frame-relay end-to-end keepalive interface S0/0.41

End-to-end Keepalive Statistics for Interface Serial0/0.41 (Frame Relay DTE)

DLCI = 401, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP)

SEND SIDE STATISTICS

Send Sequence Number: 11,   Receive Sequence Number: 12
Configured Event Window: 5,   Configured Error Threshold: 3
Total Observed Events: 14,   Total Observed Errors: 0
Monitored Events: 5,   Monitored Errors: 0
Successive Successes: 5,   End-to-end VC Status: UP

```
RECEIVE SIDE STATISTICS

Send Sequence Number: 11,        Receive Sequence Number: 10
Configured Event Window: 5,      Configured Error Threshold: 3
Total Observed Events: 13,       Total Observed Errors: 0
Monitored Events: 5,             Monitored Errors: 0
Successive Successes: 5,         End-to-end VC Status: UP
```

Task 5

Erase the startup config and reload the routers before proceeding to the next lab

Lab 7 – Tricky Frame-relay configuration

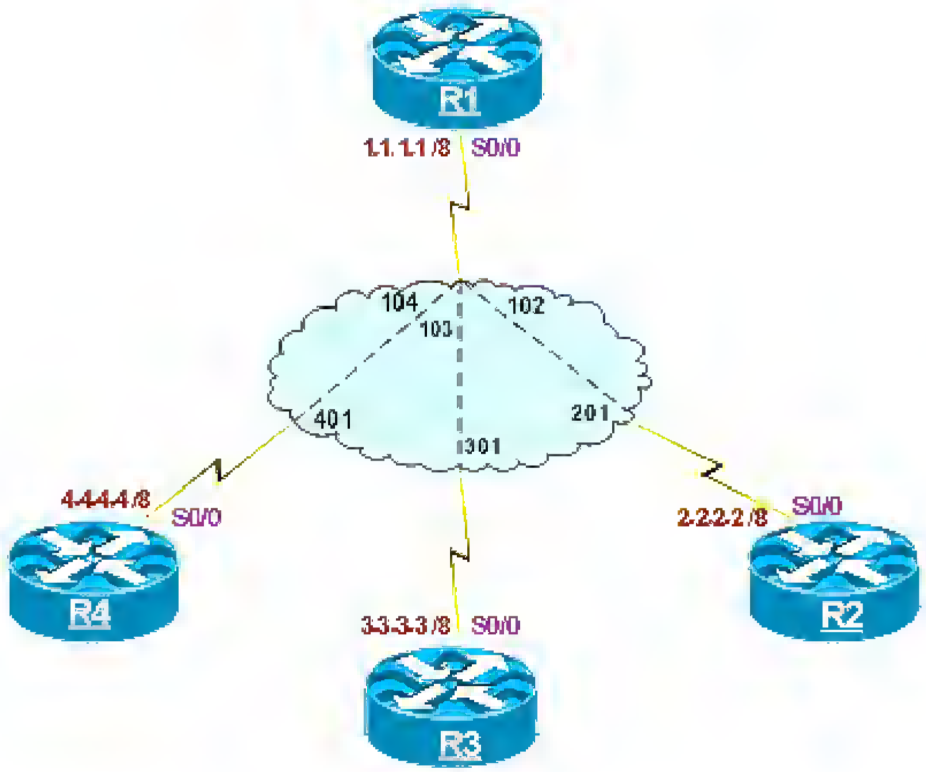

IP addressing and DLCI information Chart:

Routers	IP address	Local DLCI	Connecting to:
R1's Loopback 0 interface	1.1.1.1 /8		
R1's Frame-relay interface	Ip unnumbered Lo0	102	R2
	Ip unnumbered Lo0	103	R3
	Ip unnumbered Lo0	104	R4
R1's Loopback 0 interface	**2.2.2.2 /8**		
R2's Frame-relay interface	**Ip unnumbered Lo0**	**201**	**R1**
R3's Loopback 0 interface	3.3.3.3 /8		
R3's Frame-relay interface	Ip unnumbered Lo0	301	R1
R4's Loopback 0 interface	**4.4.4.4 /8**		
R4's Frame-relay interface	**Ip unnumbered Lo0**	**401**	**R1**

Task 1

Configure the routers in a hub and spoke manner using the IP addressing in the above chart.

The hub router (R1): This router should use DLCIs 102, 103 and 104 for it's connection to R2, R3 and R4 respectively. This router should be configured in a multipoint manner.

The spokes, R2, R3 and R4: DLCIs 201, 301 and 401 should be used by R2, R3 and R4 respectively for their frame-relay connection to R1 (The hub).

Ensure that these routers have full reachability to every Loopback interface, this should include their own. You should NOT use "**Frame-relay map**", and/or static/dynamic routing to accomplish this task.

None of the routers should be configured with sub-interface/s.

On R1

R1(config)#Int S0/0
R1(config-if)#Encap frame-relay
R1(config-if)#**Frame-relay interface-dlci 102 ppp virtual-template 1**
R1(config-if)#**Frame-relay interface-dlci 103 ppp virtual-template 1**
R1(config-if)#**Frame-relay interface-dlci 104 ppp virtual-template 1**

R1(config)#Int Virtual-template 1
R1(config-if)#Ip unnumbered lo0

R1(config)#Int lo0
R1(config-if)#Ip address 1.1.1.1 255.0.0.0

On R2

R2(config)#Int S0/0
R2(config-if)#Encap frame-relay
R2(config-if)#**Frame-relay interface-dlci 201 ppp virtual-template 2**

R2(config)#Int Virtual-template 2
R2(config-if)#Ip unnumbered lo0

R2(config)#Int lo0
R2(config-if)#Ip address 2.2.2.2 255.0.0.0

On R3

R3(config)#Int S0/0
R3(config-if)#Encap frame-relay
R3(config-if)#**Frame-relay interface-dlci 301 ppp virtual-template 3**

R3(config)#Int Virtual-template 3
R2)config-if)#Ip unnumbered lo0

R3(config)#Int lo0
R3(config-if)#Ip address 3.3.3.3 255.0.0.0

On R4

R4(config)#Int S0/0
R4(config-if)#Encap frame-relay
R4(config-if)#**Frame-relay interface-dlci 401 ppp virtual-template 4**

R4(config)#Int Virtual-template 4
R4(config-if)#Ip unnumbered lo0

R4(config)#Int lo0
R4(config-if)#Ip address 4.4.4.4 255.0.0.0

To verify and test connectivity between the hub and it's attached spokes:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     2.0.0.0/32 is subnetted, 1 subnets
C       2.2.2.2 is directly connected, Virtual-Access1
```

3.0.0.0/32 is subnetted, 1 subnets
C   3.3.3.3 is directly connected, Virtual-Access2
4.0.0.0/32 is subnetted, 1 subnets
C   4.4.4.4 is directly connected, Virtual-Access3

Note when PPP is configured, in the last step of PPP connection, IPCP creates a host route for the router's interface that is connected to your local router. This behavior can be disabled using the "no peer neighbor-route" command. Note because of this behavior in PPP, R1 should have connectivity to every spoke, as follows:

On R1

R1#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

R1#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B – BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 1 subnets
C       1.1.1.1 is directly connected, Virtual-Access1
C    2.0.0.0/8 is directly connected, Loopback0
```

Note R2 has reachability to R1 but NOT to any of the spokes

R2#Ping 1.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

R2#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

R2#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 1 subnets
C       1.1.1.1 is directly connected, Virtual-Access1
C    3.0.0.0/8 is directly connected, Loopback0
```

Note R3 has reachability to R1 but NOT to any of the spokes

R3#Ping 1.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

R3#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

R3#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R4

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

1.0.0.0/32 is subnetted, 1 subnets
C 1.1.1.1 is directly connected, Virtual-Access1
C 2.0.0.0/8 is directly connected, Loopback0

Note R4 has reachability to R1 but NOT to any of the spokes

R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

PBR can be configured to provide reachability between the spokes as follows:

On R2, R3 and R4

(config)#**Ip local policy route-map TST**

(config-route-map)#**Route-map TST permit 10**
(config-route-map)#**Set ip next-hop** 1.1.1.1

(config-route-map)#**Route-map TST permit 20**

To test the configuration:

On R2

R2#Debug ip policy

R2#Ping 3.3.3.3 source 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/117/120 ms

IP: s=2.2.2.2 (local), d=3.3.3.3, len 100, policy match
IP: route map TST, item 10, permit
IP: s=2.2.2.2 (local), d=3.3.3.3 (Virtual-Access2), len 100, policy routed
```

R2#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```

On R3

R3#Ping 1.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

R3#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```

R3#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms

On R4

R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/116 ms

R4#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Task 2

Erase the startup config and reload the routers before proceeding to the next lab

Lab 8 – Frame-relay Multilinking

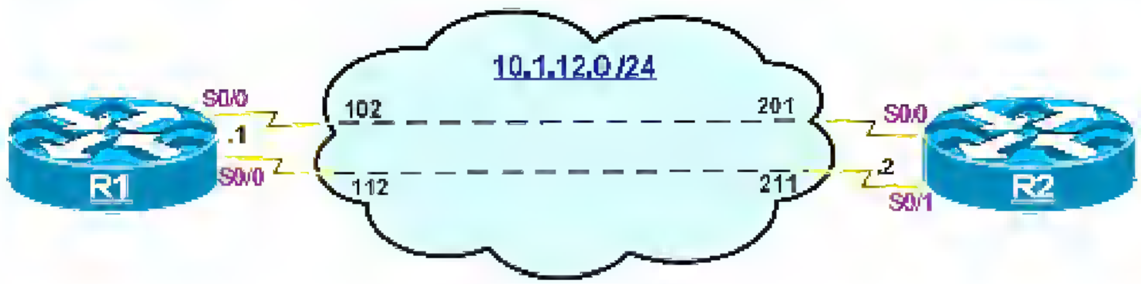

Task 1

Configure the frame-relay connections between R1 and R2 in a point-to-point manner using the DLCIs and interfaces in the diagram. Configure R1 and R2 using 10.1.12.1 /24 and 10.1.12.2 /24 IP addresses respectively. Ensure the these links appear as one and have authentication capability.

Note the task does NOT specifically ask for PPP Multilink to be configured, but since the task asks for each router to have a single IP address and it states that the links should appear as one with authentication capability, that should be enough to indicate the PPP Multilink configuration.
Most of the time there is only a single connection between two routers, but there are situations where you may need to have multiple layer one connections between the two routers, one reason could be to increase the size of the pipe between the two routers. The point of Multilink PPP is to take multiple PPP links and "bond" them together to act as a single PPP link. These PPP links that are being bonded could be an ISDN BRI circuit, T1 circuits, or other types of PPP circuits as long as they are from the same provider.

On R1

The following command creates a logical multilink group, in the following configuration the multilink group is assigned a value of 12, but the range is 1 – 2.14 Billion.
R1(config)#**int Multilink 12**

An IP address is assigned to this logical interface, as follows:

R1(config-if)#**ip addr 10.1.12.1 255.255.255.0**

To verify the configuration:

On R1

R1#Show run int multilink 12

```
Building configuration...
Current configuration : 89 bytes
!
interface Multilink12
 ip address 10.1.12.1 255.255.255.0
 ppp multilink
 ppp multilink group 12
```

The "PPP Multilink" command enables the interface to support MLP (Multilink Point-to-point Protocol) and the "PPP multilink group 12" command identifies the Multilink group that will later be assigned to two or more interfaces that will restrict them to joining only the designated multilink-group.

The following command creates a virtual-template interface and assigns the multilink group 12 to this logical interface.

R1(config)#**inter virtual-template 12**
R1(config-if)#**ppp multilink group 12**

Finally, the virtual-template 12 is assigned to the DLCIs:

R1(config-if)#**int s0/0**
R1(config-if)#**encap frame**

R1(config-if)#**int s0/0.12 Multipoint**
R1(config-subif)#**frame-relay interface-dlci 102 ppp virtual-template 12**
R1(config-subif)#**frame-relay interface-dlci 112 ppp virtual-template 12**

R1(config)#**int s0/0**
R1(config-if)#**No shut**

To verify the configuration:

On R1

R1#Show ppp multilink

No active bundles
Multilink12 (inactive)
  Member links: 2
    Vi3 (inactive)
    Vt12 (inactive)

Note the output of the above command shows that there are no active bundles; this is because PPP Multilinking must be configured on both end points before its activated.

On R2

R2(config)#**int multilink 21**
R2(config-if)#**ip addr 10.1.12.2 255.255.255.0**

R2(config)#**int virtual-template 21**
R2(config-if)#**ppp multilink group 21**

R2(config)#**int s0/0**
R2(config-if)#**encap frame-relay**
R2(config-if)#**frame-relay interface-dlci 201 ppp virtual-Template 21**
R2(config-if)#**no shut**

R2(config-if)#**int s0/1**
R2(config-if)#**encap frame-relay**
R2(config-if)#**frame-relay interface-dlci 211 ppp virtual-Template 21**
R2(config-if)#**no shut**

Note on R2 the virtual-template is assigned to two different physical interfaces, and frame-relay is configured directly under the physical interfaces. This is done intentionally to show the different implementations of this configuration.

To verify the configuration:

On R1

Note the Multilink12 logical interface is now up, this is because both routers/end points are configured with PPP multilink.

R1#Show ppp multilink

Multilink12, bundle name is R2

  Endpoint discriminator is R2

```
Bundle up for 00:16:04, total bandwidth 200000, load 1/255
  Receive buffer limit 24000 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 0 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x24 received sequence, 0x24 sent sequence
  Member links: 2 active, 1 inactive (max not set, min not set)
    Vi1, since 00:16:04
    Vi2, since 00:16:03
    Vt12 (inactive)
No inactive multilink interfaces
```

To verify the configuration:

On R2

R2#Show ppp multilink

```
Multilink21, bundle name is R1
  Endpoint discriminator is R1
  Bundle up for 00:18:19, total bandwidth 200000, load 1/255
  Receive buffer limit 24000 bytes, frag timeout 1000 ms
    0/0 fragments/bytes in reassembly list
    0 lost fragments, 0 reordered
    0/0 discarded fragments/bytes, 0 lost received
    0x28 received sequence, 0x28 sent sequence
  Member links: 2 active, 1 inactive (max not set, min not set)
    Vi1, since 00:18:19
    Vi2, since 00:18:19
    Vt21 (inactive)
No inactive multilink interfaces
```

To test the configuration:

On R1

R1#Ping 10.1.12.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms
```

R1#Show ip route | b Gateway

```
Gateway of last resort is not set
     10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.1.12.2/32 is directly connected, Multilink12
C       10.1.12.0/24 is directly connected, Multilink12
```

Note the host route is installed because of PPP implementation.

Task 2

Configure CHAP authentication between the two routers. Use "**Cisco**" as the password.

On R1

R1(config)#**username R2 password Cisco**

R1(config)#**int virtual-template 12**
R1(config-if)#**ppp authentication chap**

On R2

R2(config)#**username R1 password Cisco**

R2(config)#**int virtual-template 21**
R2(config-if)#**ppp authentication chap**

Note the authentication is configured under the virtual-template interface.

To verify the configuration:

On R1

R2#Show ppp multilink

```
Multilink12, bundle name is R2
  Username is R2
  Endpoint discriminator is R2
  Bundle up for 00:00:28, total bandwidth 200000, load 1/255
  Receive buffer limit 24000 bytes, frag timeout 1000 ms

    0/0 fragments/bytes in reassembly list
    8 lost fragments, 10 reordered
```

```
    4/350 discarded fragments/bytes, 2 lost received
    0x12 received sequence, 0x2 sent sequence
  Member links: 2 active, 1 inactive (max not set, min not set)
    Vi1, since 00:00:29
    Vi2, since 00:00:18
    Vt12 (inactive)
```

Note this line is added and it indicates that authentication is configured.

To test the configuration:

On R2

R2#Ping 10.1.12.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/55/56 ms
```

Task 3

Erase the startup configuration and reload the routers before proceeding to the next section.

Lab 9 – Back-to-Back Frame-relay connection

Lab Setup:

- No Setup is necessary; this lab is configured on the serial interface of R1 that is directly connected to R3 without the presence of a frame-relay switch.

IP addressing:

Router	Interface / IP address	DLCI assignment
R1	S0/1 = 200.1.1.1 /24	113
R3	S0/1 = 200.1.1.3 /24	113

Task 1

Configure Frame-relay between R1 and R3, you should use the IP address, interface and the DLCIs provided in the IP Addressing table above.

In this scenario we do not have a frame-relay switch connecting the routers; these routers are connected back to back using a DTE ←→ DCE serial cable. The router that is connected to the DCE side should provide the clocking using the "Clock rate" interface configuration command, the DCE side can be

determined using the "Show controller S 0/1" command as follows:

R1#Sh controller S 0/1

CD2430 Slot 1, Port 0, Controller 0, Channel 0, Revision 19
Channel mode is synchronous serial
idb 0x84E4BAB8, buffer size 1524, V.35 DCE cable
(The rest of the output is omitted)

In this case since the frame-relay switch does NOT exist, the LMIs should be disabled using the "No Keepalive" interface configuration command, and the frame-relay mapping should be done statically.
When configuring the Frame-relay mapping, the DLCIs should be identical on both ends.

On R1

R1(config)#interface Serial0/1
R1(config-if)#ip address 200.1.1.1 255.255.255.0
R1(config-if)# encapsulation frame-relay
R1(config-if)# **no keepalive**
R1(config-if)# clock rate 64000
R1(config-if)# **frame-relay map ip 200.1.1.3** 113

On R3

R3(config)#interface Serial0/1
R3(config-if)#ip address 200.1.1.3 255.255.255.0
R3(config-if)# encapsulation frame-relay
R3(config-if)# **no keepalive**
R3(config-if)# **frame-relay map ip 200.1.1.1** 113

To verify & test the configuration:

On R1

R1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

R1#Show frame-relay lmi

R1#

Note there are no LMIs, because they are disabled.

R1#Show frame-relay pvc

PVC Statistics for interface Serial0/1 (Frame Relay DTE)

	Active	Inactive	Deleted	Static
Local	0	0	0	1
Switched	0	0	0	0
Unused	0	0	0	0

DLCI = 113, DLCI USAGE = LOCAL, PVC STATUS = STATIC, INTERFACE = Serial0/1

```
input pkts 5          output pkts 5          in bytes 520
out bytes 520         dropped pkts 0         in pkts dropped 0
out pkts dropped 0             out bytes dropped 0
in FECN pkts 0        in BECN pkts 0         out FECN pkts 0
out BECN pkts 0       in DE pkts 0           out DE pkts 0
out bcast pkts 0      out bcast bytes 0
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
pvc create time 00:29:24, last time pvc status changed 00:29:24
```

R1#Show frame-relay map

Serial0/1 (up): **ip 200.1.1.3 dlci 113**(0x71,0x1c10), **static**,
CISCO

Task 2

Reconfigure the routers such that R1 uses DLCI 103 to send and DLCI 301 to receive packets, whereas, R3 should use DLCI 301 to send and DLCI 103 to receive packets. You should configure interface S0/1 to accomplish this task.

In this configuration, we are asked to configure these routers using different DLCIs, 103 connecting R1 to R3 and 301 connecting R3 to R1.

On R1

R1(config)#interface Serial0/1

R1(config-if)#ip address 200.1.1.1 255.255.255.0
R1(config-if)# encapsulation frame-relay
R1(config-if)# no keepalive
R1(config-if)# clock rate 64000

The following command removes the frame-relay mapping that was configured in the previous task and adds the new mapping:

R1(config-if)#**NO frame-relay map ip 200.1.1.3 113**
R1(config-if)#**frame-relay map ip 200.1.1.3 103**

On R3

R3(config)#interface Serial0/1
R3(config-if)#ip address 200.1.1.3 255.255.255.0
R3(config-if)#encapsulation frame-relay
R3(config-if)#no keepalive

R3(config-if)#**NO frame-relay map ip 200.1.1.1 131**
R3(config-if)#**frame-relay map ip 200.1.1.1 301**

To verify and test the configuration:

On Both Routers:

#Debug ip packet
#Debug Frame-relay packet

On R1

R1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note the ping is NOT successful and the following messages on R3 will reveal the reason:

Serial0/1: FR invalid/unexpected pak received on DLCI 103
Serial0/1: FR invalid/unexpected pak received on DLCI 103
Serial0/1: FR invalid/unexpected pak received on DLCI 103
Serial0/1: FR invalid/unexpected pak received on DLCI 103

Serial0/1: FR invalid/unexpected pak received on DLCI 103

Note the above errors are received on R3 because the DLCIs don't match, R1 does not know about DLCI 103. Remember that they are connected directly.
To fix this problem, R3 can be configured to receive data on DLCI 103 and send on DLCI 301, as follows:

On R3

R3(config)#int S0/1
R3(config-if)#frame-relay interface-dlci 103

To verify and test the configuration:

On R3

R3#Debug frame-relay packet

On R1

R1#Ping 200.1.1.3 repeat 4

On R3

Serial0/1(i): dlci 103(0x1871), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 301(0x48D1), pkt type 0x800(IP), datagramsize 104

Serial0/1(i): dlci 103(0x1871), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 301(0x48D1), pkt type 0x800(IP), datagramsize 104

Serial0/1(i): dlci 103(0x1871), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 301(0x48D1), pkt type 0x800(IP), datagramsize 104

Serial0/1(i): dlci 103(0x1871), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 301(0x48D1), pkt type 0x800(IP), datagramsize 104

Note the incoming traffic uses DLCI 103, whereas, the outgoing traffic uses DLCI 301.

To test the configuration:

On R1

R1#Debug Frame-relay Packet

On R3

R3#Ping 200.1.1.1 repeat 4

On R1

Serial0/1: FR invalid/unexpected pak received on DLCI 301
Serial0/1: FR invalid/unexpected pak received on DLCI 301
Serial0/1: FR invalid/unexpected pak received on DLCI 301
Serial0/1: FR invalid/unexpected pak received on DLCI 301

Note the same problem, the traffic comes in on DLCI 301 and the local router is NOT aware of this DLCI. To fix the problem:

R1(config)#int S0/1
R1(config-if)#**frame-relay interface-dlci 301**

To verify and test the configuration:

On R3

R3#Ping 200.1.1.1 repeat 4

On R1

Serial0/1(i): dlci 301(0x48D1), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 103(0x1871), pkt type 0x800(IP), datagramsize 104

Serial0/1(i): dlci 301(0x48D1), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 103(0x1871), pkt type 0x800(IP), datagramsize 104

Serial0/1(i): dlci 301(0x48D1), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 103(0x1871), pkt type 0x800(IP), datagramsize 104

Serial0/1(i): dlci 301(0x48D1), pkt type 0x800, datagramsize 104
Serial0/1(o): dlci 103(0x1871), pkt type 0x800(IP), datagramsize 104

R1#Show frame map

Serial0/1 (up): **ip 200.1.1.3 dlci 103(0x67,0x1870), static,**
CISCO

On R3

R3#Show frame map

Serial0/1 (up): ip 200.1.1.1 dlci 301(0x12D,0x48D0), static,
CISCO

To test connectivity:

On R1

R1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/33 ms

On R3

R3#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

Task 3

Re-configure R1 as the frame-relay switch and a router connecting to R3, whereas, R3 should be configured as a router connecting to R1 using S0/1 interface. R1 should use DLCI 103 for its connection to R3 and R3 should use DLCI 301 for its connection to R1. You should NOT disable LMIs to accomplish this task.

On R1

R1(config)#frame switching

R1(config)#int S0/1
R1(config-if)#ip addr 200.1.1.1 255.255.255.0

R1(config-if)#encap frame-relay

R1(config-if)#clock rate 64000
R1(config-if)#**frame map ip 200.1.1.3 103**
R1(config-if)#**frame interface-dlci 301**
R1(config-if)#frame-relay intf-type dce

On R3

R3(config-if)#int S0/1
R3(config-if)#ip addr 200.1.1.3 255.255.255.0
R3(config-if)#encap frame-relay
R3(config-if)#**frame map ip 200.1.1.1 301**

To verify and test the configuration:

On R1

R1#Show frame lmi | B Num

```
 Num Status Enq. Rcvd 11          Num Status msgs Sent 11
 Num Update Status Sent 0         Num St Enq. Timeouts 0
```

On R3

R3#Show frame-relay lmi | B Num

```
 Num Status Enq. Sent 18          Num Status msgs Rcvd 19
 Num Update Status Rcvd 0         Num Status Timeouts 0
 Last Full Status Req 00:00:00    Last Full Status Rcvd 00:00:00
```

R3#Show frame-relay pvc 301

```
PVC Statistics for interface Serial0/1 (Frame Relay DTE)

DLCI = 301, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0/1

  input pkts 3            output pkts 5           in bytes 102
  out bytes 520           dropped pkts 0          in pkts dropped 0
  out pkts dropped 0               out bytes dropped 0
  in FECN pkts 0          in BECN pkts 0          out FECN pkts 0
  out BECN pkts 0         in DE pkts 0            out DE pkts 0
  out bcast pkts 0        out bcast bytes 0
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
```

```
 pvc create time 00:06:03, last time pvc status changed 00:02:42

R3#Show frame-relay map

Serial0/1 (up): ip 200.1.1.1 dlci 301(0x12D,0x48D0), static,
           CISCO, status defined, active

R3#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/33 ms
```

Task 4

Erase the startup configuration and reload the routers before proceeding to the next lab.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

On Demand Routing

Lab 1 – On Demand Routing

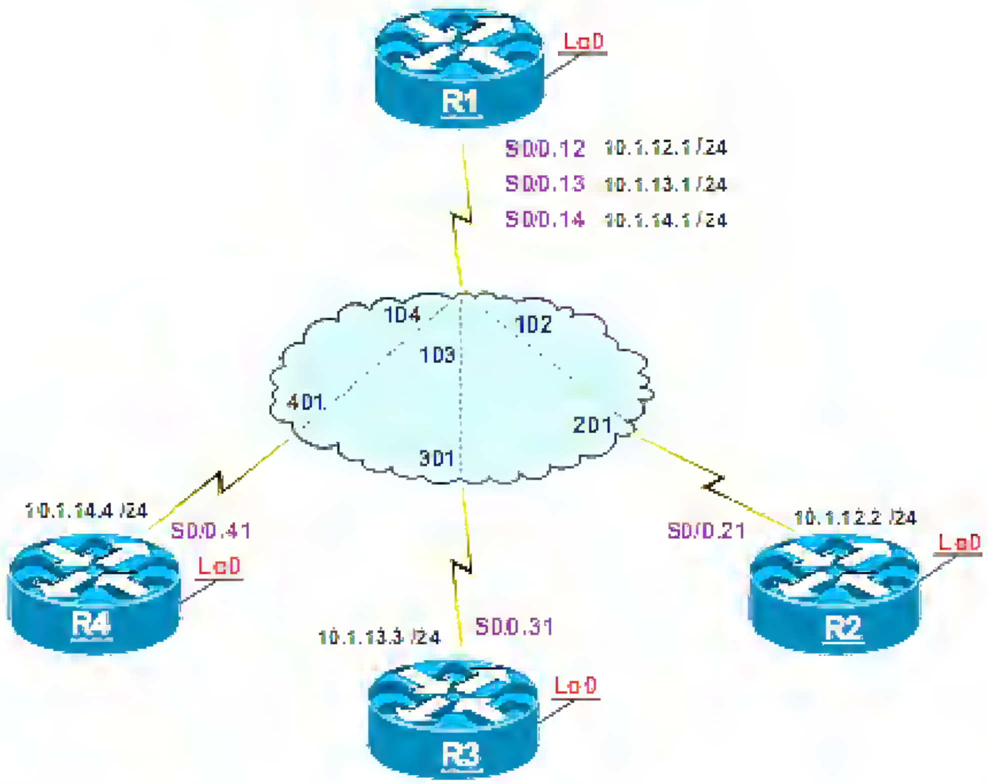

Lab Setup:

- Configure all frame-relay connections in point to point manner, with R1 as the hub and R2 - R4 as the spoke routers.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address	DLCI assignment
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24 S0/0.14 = 10.1.14.1 /24 Loopback0 = 1.1.1.1 /8	102 103 104
R2	S0/0.21 = 10.1.12.2 /24 Loopback0 = 2.2.2.2 /8	201
R3	S0/0.31 = 10.1.13.3 /24 Loopback0 = 3.3.3.3 /8	301
R4	S0/0.41 = 10.1.14.4 /24 Loopback0 = 4.4.4.4 /8	401

Task 1

Configure ODR on the appropriate router and ensure full connectivity between the routers.

On R1

R1(config)#Router odr

R1#Sh ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
o    2.0.0.0/8 [160/1] via 10.1.12.2, 00:00:20, Serial0/0.12
o    3.0.0.0/8 [160/1] via 10.1.13.3, 00:00:18, Serial0/0.13
```

```
o    4.0.0.0/8 [160/1] via 10.1.14.4, 00:00:35, Serial0/0.14
     10.0.0.0/24 is subnetted, 3 subnets
C       10.1.14.0 is directly connected, Serial0/0.14
C       10.1.13.0 is directly connected, Serial0/0.13
C       10.1.12.0 is directly connected, Serial0/0.12
```

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.1.12.1 to network 0.0.0.0

C    2.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0.21
o*   0.0.0.0/0 [160/1] via 10.1.12.1, 00:00:28, Serial0/0.21
```

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.1.13.1 to network 0.0.0.0

C    3.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
o*   0.0.0.0/0 [160/1] via 10.1.13.1, 00:00:43, Serial0/0.31
```

R4#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.1.14.1 to network 0.0.0.0

C    4.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.14.0 is directly connected, Serial0/0.41
o*   0.0.0.0/0 [160/1] via 10.1.14.1, 00:00:45, Serial0/0.41

R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/112 ms

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

ODR is a nice solution in a small hub and spoke scenario where the spokes are stub networks. ODR uses CDP as its transport. Ensure that the CDP versions match.

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

RIPv2

Lab 1 – RIPv2 and Frame-Relay

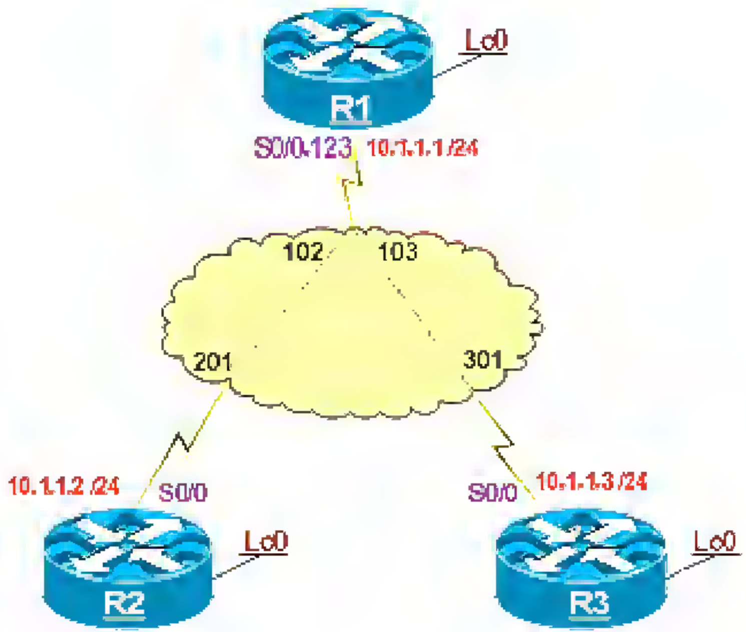

Lab Setup:

- Configure R1 as the hub and routers R2 and R3 as the spokes.
- Configure all routers in a Frame-relay Multipoint manner. DO NOT configure sub-interfaces on R2 or R3. R1 should be configured with a sub-interface in a multipoint manner. Use the broadcast keyword where necessary.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0.123 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	S0/0 = 10.1.1.2 /24 Loopback0 = 2.2.2.2 /8
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8

Task 1

Configure RIPv2 on all routers and advertise their directly connected interfaces in this routing protocol. Ensure that these routers have full NLRI to all the loopback interfaces advertised in this routing protocol.

On R1

```
R1(config-if)#router rip
R1(config-router)#ver 2
R1(config-router)#no au
R1(config-router)#netw 10.0.0.0
R1(config-router)#netw 1.0.0.0
```

On R2

```
R2(config-if)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
R2(config-router)#netw 2.0.0.0
```

On R3

```
R3(config-if)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

R3(config-router)#netw 3.0.0.0

To verify the configuration:

On R1

R1#Show ip route rip

```
R    2.0.0.0/8 [120/1] via 10.1.1.2, 00:00:07, Serial0/0.123
R    3.0.0.0/8 [120/1] via 10.1.1.3, 00:00:15, Serial0/0.123
```

On R2

R2#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:07, Serial0/0
```

On R3

R3#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:23, Serial0/0
```

Note the spoke routers do not see each others loopback interfaces; this is because of IP split horizon. If the hub router is configured in a multipoint manner using a sub-interface, then, the IP split horizon is ENABLED, whereas, if the hub is configured in a multipoint manner using the physical interface, then, the Split horizon is DISABLED, the following show commands will reveal this information:

Note the following shows the sub-interface S0/0.123 which is configured in a multipoint manner and it also reveals that IP split horizon is enabled.

On R1

R1#Show ip interface S0/0.123 | Inc Split

```
  Split horizon is enabled
```

The following command shows that S0/0 interface of R2 which is configured in a multipoint manner has its IP split horizon disabled.

On R2

R2#Sh ip int S0/0 | Inc Split

Split horizon is **disabled**

To resolve this problem:

On R1

R1(config)#int S0/0.123
R1(config-subif)#**NO ip split-horizon**

To verify the configuration:

On R1

R1#Sh ip int s0/0.123 | Inc Split

Split horizon is disabled

On R2

R2#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:12, Serial0/0
R    3.0.0.0/8 [120/2] via 10.1.1.3, 00:00:12, Serial0/0
```

On R3

R3#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:22, Serial0/0
R    2.0.0.0/8 [120/2] via 10.1.1.2, 00:00:22, Serial0/0
```

Task 2

Ensure that every router can Ping every loopback interface advertised in this routing domain.

On R2

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

On R3

R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

Note even though the prefixes advertised are in every router's routing table, R2 and R3 do NOT have reachability to each other's loopback interface:

On R2

R2#Show ip route rip

R   1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:12, Serial0/0
R   3.0.0.0/8 [120/2] via 10.1.1.3, 00:00:12, Serial0/0

On R3

Note the next hop IP address is set Based on the originating router and NOT the advertising router

R3#Show ip route rip

R   1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:22, Serial0/0
R   2.0.0.0/8 [120/2] via 10.1.1.2, 00:00:22, Serial0/0

To fix the problem:

On R2

R2(config)#Int S0/0
R2(config-if)#Frame-relay map ip 10.1.1.3 201

On R3

R3(config)#Int S0/0

R3(config-if)#Frame-relay map ip 10.1.1.2 301

To test the configuration:

On R2

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms

On R3

R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/125/168 ms

Task 3

Remove the "no **ip split-horizon**" and the "Frame-relay map" statements from R2 and R3 that was configured in the previous steps and ensure that R2 and R3 can ping each other's loopback interface. DO NOT configure static routes or reconfigure Frame-relay to accomplish this task. Ensure that the next hop IP address is NOT changed and its still the IP address of the router that is originating the prefix.

PPP's behavior is used to accomplish this task, when PPP is running on a link, the host (Peer neighbor's routes) is added to the routers routing table.

On R1

R1(config)#int S0/0.123
R1(config-subif)#**ip split-horizon**

R1(config)#int S0/0.123

```
R1(config-subif)#NO ip address
R1(config-subif)#frame interface-dlci 102 ppp virtual-template 123
R1(config-fr-dlci)#frame interface-dlci 103 ppp virtual-template 123

R1(config-subif)#Inter virtual-template 123
R1(config-if)#ip address 10.1.1.1 255.255.255.0
```

On R2

```
R2(config)#int S0/0
R2(config-if)#NO ip address
R2(config-if)#frame interface-dlci 201 ppp virtual-Template 123

R2(config-if)#int virtual-template 123
R2(config-if)#ip addr 10.1.1.2 255.255.255.0
```

On R3

```
R3(config)#int S0/0
R3(config-if)#NO ip address
R3(config-if)#frame interface-dlci 301 ppp virtual-Template 123

R3(config-if)#int virtual-template 123
R3(config-if)#ip addr 10.1.1.3 255.255.255.0
```

To verify the configuration:

On R1

```
R1#Show ip route rip

R    2.0.0.0/8 [120/1] via 10.1.1.2, 00:00:08, Virtual-Access1
R    3.0.0.0/8 [120/1] via 10.1.1.3, 00:00:20, Virtual-Access2
```

On R2

Note the next hop IP address is NOT changed

```
R2#Show ip route rip
```

The peer neighbor route added

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:26, Virtual-Access1
R    3.0.0.0/8 [120/2] via 10.1.1.3, 00:00:26, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R       10.1.1.3/32 [120/1] via 10.1.1.1, 00:00:26, Virtual-Access1
```

On R3

R3#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:06, Virtual-Access1
R    2.0.0.0/8 [120/2] via 10.1.1.2, 00:00:06, Virtual-Access1
     10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R       10.1.1.2/32 [120/1] via 10.1.1.1, 00:00:06, Virtual-Access1
```

To test the configuration:

On R2

R2#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/116/117 ms
```

On R3

R3#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```

Task 4

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 2 – RIPv2 Authentication

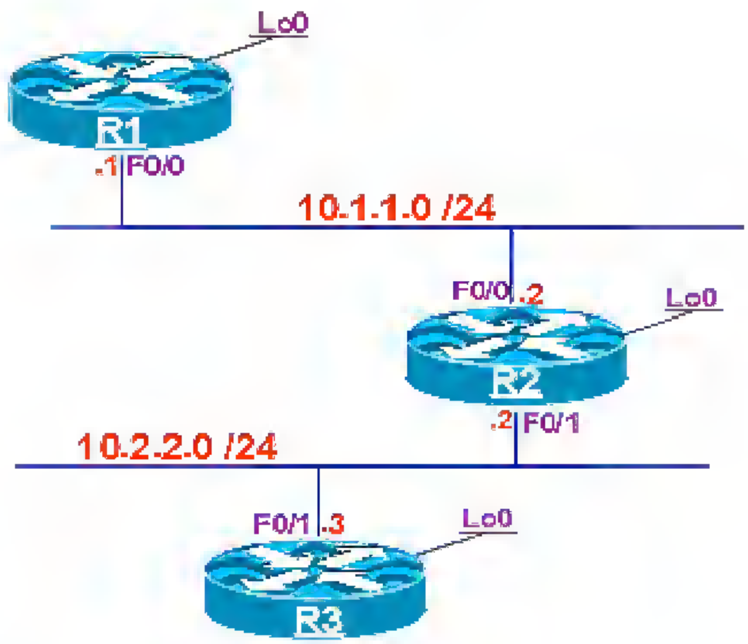

Lab Setup:

- Configure the F0/0 interface of R1 and R2 in VLAN 12.
- Configure the F0/1 interface of R2 and R3 in VLAN 23
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	F0/0 = 10.1.1.2 /24 F0/1 = 10.2.2.2 /24 Loopback0 = 2.2.2.2 /8
R3	F0/1 = 10.2.2.3 /24 Loopback0 = 3.3.3.3 /8

Task 1

Configure RIPv2 on R1 and R2 and advertise their directly connected networks in this routing protocol. You should disable auto summarization when configuring RIP.

On R1

```
R1(config-if)#router rip
R1(config-router)#ver 2
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 10.0.0.0
```

On R2

```
R2(config-if)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
R2(config-router)#netw 2.0.0.0
```

To verify the configuration:

On R1

R1#Show ip route rip

```
R    2.0.0.0/8 [120/1] via 10.1.1.2, 00:00:15, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
R       10.2.2.0 [120/1] via 10.1.1.2, 00:00:27, FastEthernet0/0
```

On R2

R2#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:16, FastEthernet0/0
```

Task 2

Configure clear text RIPv2 authentication between R1 and R2. Use "cisco" as the password for this authentication.

On Both Routers

(config)#Key chain TST
(config-keychain)#key 1
(config-keychain-key)#**key-string cisco**

(config-keychain-key)#int f0/0
(config-if)#**ip rip authentication key-chain TST**

To verify the configuration:

On R1

R1#Show ip route rip

```
R    2.0.0.0/8 [120/1] via 10.1.1.2, 00:00:01, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
R       10.2.2.0 [120/1] via 10.1.1.2, 00:00:01, FastEthernet0/0
```

Note in RIPv2 there are two types of authentication, Clear text and MD5. In RIPv2 clear text authentication the key numbers do not need to match, meaning that R1 can use one key number and R2 can use a totally different one.

Task 3

Configure RIPv2 MD5 authentication between R2 and R3, these routers should use "**cisco23**" as the password for this authentication.

On R3

R3(config-subif)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
R3(config-router)#netw 3.0.0.0

On Routers R2 and R3

(config)#key chain TST23
(config-keychain)#key 1
(config-keychain-key)#key-string cisco23

(config)#int f0/1
(config-if)#**ip rip authentication key-chain** TST23
(config-if)#**ip rip authentication mode md5**

To verify the configuration:

On R2

R2#Sh ip route rip

R   1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:20, FastEthernet0/0
R   3.0.0.0/8 [120/1] via 10.2.2.3, 00:00:16, FastEthernet0/1

Note when configuring MD5 authentication, the passwords and the key numbers MUST match on both routers, or else the routers will not exchange routing updates.

Task 4

Configure R2 such that it receives all routes from R3, whereas, R3 ignores v2 packet/s from R2 (10.2.2.2). DO NOT use any filtering, offset-list, route-maps or passive-interface to accomplish this task.

In this case the behavior of MD5 authentication is used to accomplish this task, in RIPv2 MD5 authentication if the key numbers do NOT match, the router with a higher key number will receive all the routes and it will populate the received routes in its routing table, whereas, the router that has a lower key number will totally ignore all routes received from the other router.

On R2

R2(config)#**No key chain TST23**

R2(config)#key chain TST23
R2(config-keychain)#**key 2**

The key number of R2 is 2, whereas, the key number R3 is 1.

R2(config-keychain-key)#key-string cisco23

On R3

R3#Show key chain TST23

Key-chain TST23:
key 2 -- text "cisco23"
accept lifetime (always valid) - (always valid) [valid now]
send lifetime (always valid) - (always valid) [valid now]

To test the configuration:

On R3

R3#Show ip route rip
R3#

On R2

R2#Show ip route rip

R 1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:18, FastEthernet0/0
R 3.0.0.0/8 [120/1] via 10.2.2.3, 00:00:27, FastEthernet0/1

Note R2 will receive and process the routes, whereas, R3 will reject the routes because the key numbers do not match and its key number is lower than R2's key number.

Task 5

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 3 – Advanced RIPv2 Mock Lab

Logical Topology

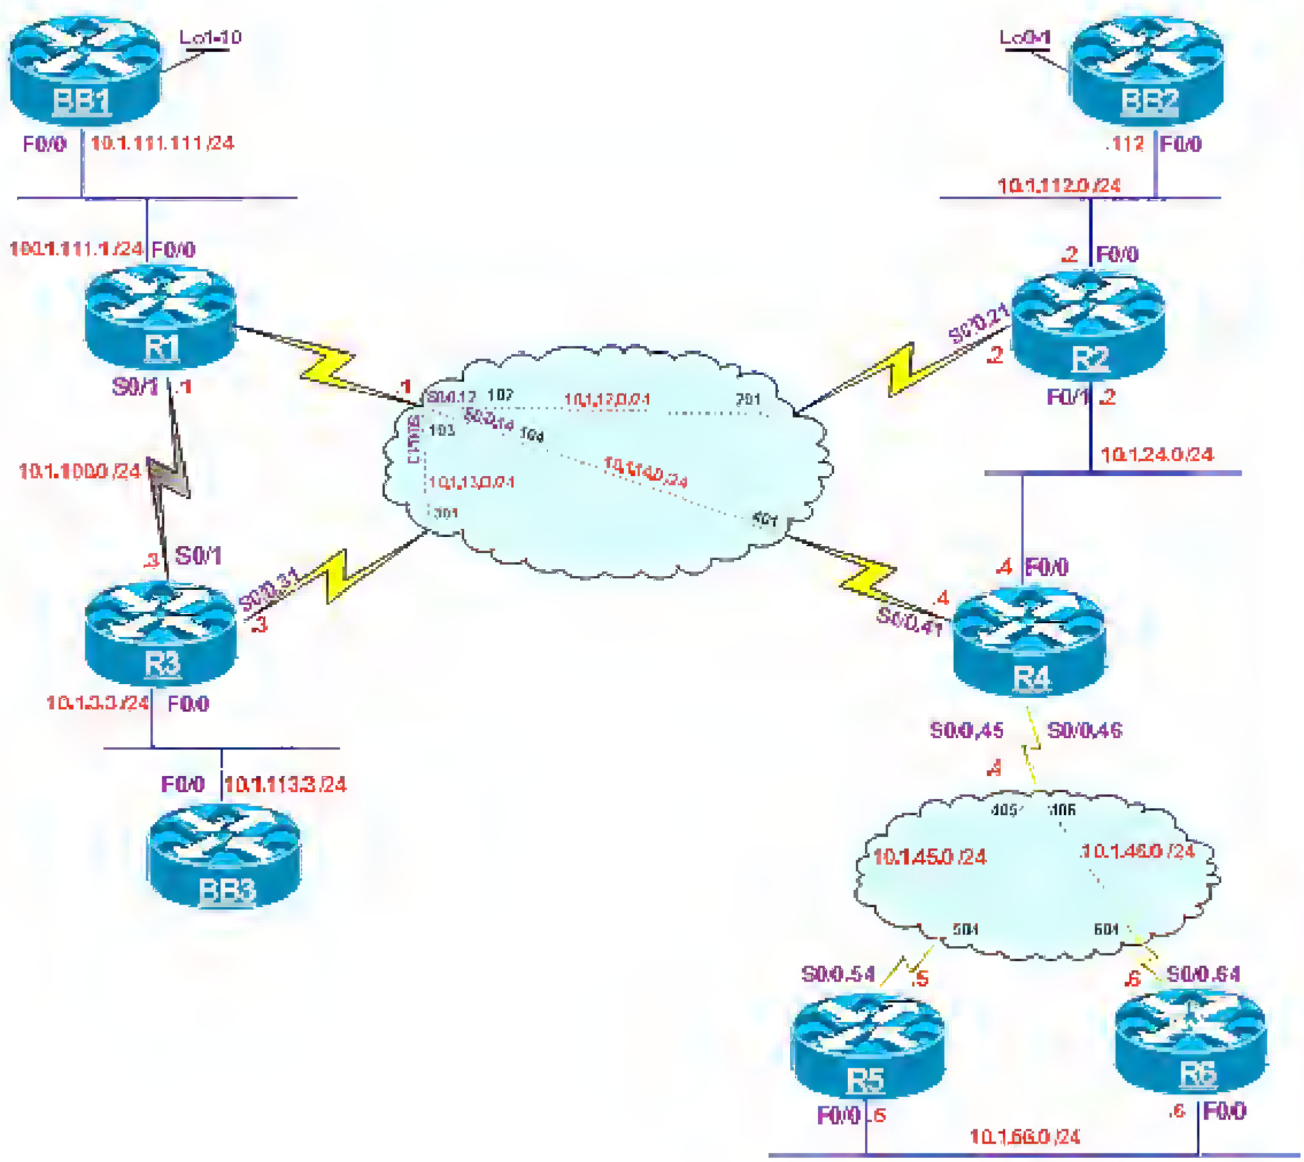

IP Addressing chart:

Router	Interface	Connecting to:	IP Address
R1	S0/0.12	R2	10.1.12.1 /24
	S0/0.13	R3	10.1.13.1 /24
	S0/0.14	R4	10.1.14.1 /24
	S0/1	R3	10.1.100.1 /24
	F0/0	BB1	100.1.111.1 /24
R2	S0/0.21	R1	10.1.12.2 /24
	F0/0	BB2	10.1.112.2 /24
	F0/1	R4	10.1.24.2 /24
R3	S0/0.31	R1	10.1.13.3 /24
	S0/1	R1	10.1.100.3 /24
	F0/0		10.1.3.3 /24
R4	S0/0.41	R1	10.1.14.4 /24
	S0/0.45	R5	10.1.45.4 /24
	S0/0.46	R6	10.1.46.4 /24
	F0/0	R2	10.1.24.4 /24
R5	S0/0.54	R4	10.1.45.5 /24
	F0/0	R6	10.1.56.5 /24
R6	S0/0.64	R4	10.1.46.6 /24
	F0/0	R5	10.1.56.6 /24
BB1	F0/0	R1	10.1.111.111 /24
	Lo1		101.0.0.111 /8
	Lo2		102.0.0.111 /9
	Lo3		103.0.0.111 /10
	Lo4		104.0.0.111 /11
	Lo5		105.0.0.111 /14
	Lo6		106.1.1.33 /27
	Lo7		107.1.1.111 /25
	Lo8		108.1.1.65 /26
	Lo9		109.1.4.111 /22
	Lo10		110.1.1.17 /28
BB2	F0/0	R2	10.1.112.112 /24
	Lo0		112.1.1.1 /24
	Lo1		112.2.2.2 /24
BB3	E0/0		10.1.113.3 /24

Lab Setup:

VLANs:

- BB1 and R1's F0/0 interface should be configured in VLAN 111
- BB2 and R2's F0/0 should be configured in VLAN 112
- R2's F0/1 and R4's F0/0 interface should be configured in VLAN 24
- BB3's E0/0 should be configured in VLAN 113.
- R3's F0/0 interface should be configured in VLAN 3
- R5 and R6 should have their F0/0 interface in VLAN 56.

Frame-relay:

- R4 should be configured with two sub-interfaces in a point-to-point manner, one connecting R4 to R5 and the second one connecting R4 to R6.
- R5 and R6 should each be configured with a single point-to-point sub-interface connection to R4.
- The frame-relay connection between R1, R2, R3 and R4 should be configured in a hub and spoke manner as follows:

  R1 should be configured with three point-to-point sub-interfaces connecting it to routers R2, R3 and R4.

  Routers R2, R3 and R4 should each be configured with a point-to-point frame-relay connection to R1.

Trunk connection between the switches:

SW-1 and SW-2 should be connected to each other via ports F0/19 and F0/20 forming an ISL trunk.

Task 1

Configure RIPv2 on the routers and advertise their directly connected interfaces in this routing domain.

On All Routers

```
(config)#Router rip
(config-router)#No au
(config-router)#Ver 2

(config-router)#Network 10.0.0.0
```

On R1

```
R1(config)#Router rip
R1(config-router)#Network 100.0.0.0
```

On R6

R6#Show ip route rip

```
     100.0.0.0/24 is subnetted, 1 subnets
R       100.1.111.0 [120/2] via 10.1.46.4, 00:00:02, Serial0/0.64
     10.0.0.0/24 is subnetted, 9 subnets
R       10.1.14.0 [120/1] via 10.1.46.4, 00:00:02, Serial0/0.64
R       10.1.13.0 [120/2] via 10.1.46.4, 00:00:02, Serial0/0.64
R       10.1.12.0 [120/2] via 10.1.46.4, 00:00:02, Serial0/0.64
R       10.1.3.0 [120/3] via 10.1.46.4, 00:00:02, Serial0/0.64
R       10.1.24.0 [120/1] via 10.1.46.4, 00:00:02, Serial0/0.64
R       10.1.45.0 [120/1] via 10.1.56.5, 00:00:01, FastEthernet0/0
                  [120/1] via 10.1.46.4, 00:00:02, Serial0/0.64
R       10.1.112.0 [120/2] via 10.1.46.4, 00:00:02, Serial0/0.64
```

Task 2

Set the RIPv2 timers on all routers to be twice as much as the default value for update, invalidation timer, hold down, and flush timer.

To find out the default parameters:

On R1

R1#Show ip proto | Inc Send|Invalid

Sending updates every 30 seconds, next due in 23 seconds
Invalid after 180 seconds, hold down 180, flushed after 240

On All Routers

(config)#Router rip
Timers Basic ?
<0-4294967295> Interval between updates
R1(config-router)#timers basic 60 ?
<1-4294967295> Invalid

R1(config-router)#timers basic 60 360 ?
<0-4294967295> Holddown

R1(config-router)#timers basic 60 360 360 ?

<1-4294967295> Flush

R1(config-router)#timers basic 60 360 360 480

On All Routers

(config-router)#Timers basic 60 360 360 480

Task 3

In order to avoid collisions, R6 should delay a regular periodic updates by up to 100 milliseconds.

On R6

R6(config-router)#Timers basic 60 360 360 480 100 This is the sleep timer

Task 4

R5 and R6 should suppress a flash update if the regular update is due in 10 seconds or less.

On R5 and R6

(config)#router rip
(config-router)#**flash-update-threshold 10**

The Flash-update-threshold command suppresses flash updates when the arrival of a regularly scheduled periodic update matches, or is less than the number of seconds that is configured, in this case 10 seconds. The range is 0 – 30 seconds.
The above configuration configures both routers to suppress a flash update, if the regular periodic update is due in 10 seconds or less.

To verify the configuration:

Rx#Show ip protocols

Routing Protocol is "rip"

  Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
  Sending updates every 60 seconds, next due in 35 seconds
  Flash update is suppressed when next update due within 10 seconds
  Invalid after 360 seconds, hold down 360, flushed after 480
  (The rest of the output is omitted)

Task 5

The link between R4 ←→ R5, R4 ←→ R6 should use authentication when exchanging routing updates, the password for this authentication should be set to "cisco", these routers should use the strongest authentication method available in RIPv2.

On R4, R5 and R6

(config)#key chain TST
(config-keychain)#key 1
(config-keychain-key)#key-string cisco

On R4

```
R4(config)#int S0/0.45
R4(config-if)#ip rip authentication key-chain TST
R4(config-if)#ip rip authentication mode md5

R4(config)#int S0/0.46
R4(config-if)#ip rip authentication key-chain TST
R4(config-if)#ip rip authentication mode md5
```

On R5

```
R5(config)#int S0/0.54
R5(config-if)#ip rip authentication key-chain TST
R5(config-if)#ip rip authentication mode md5
```

On R6

```
R6(config)#int S0/0.64
R6(config-if)#ip rip authentication key-chain TST
R6(config-if)#ip rip authentication mode md5
```

To verify the configuration:

On R4

```
R4#Show ip protocols | Inc Interface|TST

  Interface          Send  Recv  Triggered RIP  Key-chain
  Serial0/0.45       2     2                    TST
  Serial0/0.46       2     2                    TST
```

Task 6

R1 is configured with RIPv2 and it's advertising its directly connected networks. Ensure that R1 receives 10 routes from BB1. DO NOT configure tunnel, secondary IP addressing for this task. Ensure that R1 has reachability to all the networks advertised by BB1; you are allowed one static route to accomplish this task.

On R1

R1(config)#router rip
R1(config-router)#no **validate-update-source**

RIP and IGRP are the ONLY two routing protocols that validate the source IP address of incoming updates.
Before RIP and/or IGRP routing protocols accept routes from a given neighbor, they want to make sure that the source IP address of the advertising router is from the same IP address space as the link that the two routers are connected to. If the routers that have to exchange routing information are from different IP address spaces, then the source validation should be negated using "no validate-update-source"

To Verify the configuration:

On R1

R1#Show ip route rip | Inc 10.1.111.111

```
R       102.0.0.0 [120/1] via 10.1.111.111, 00:00:12
R       103.0.0.0 [120/1] via 10.1.111.111, 00:00:12
R    101.0.0.0/8 [120/1] via 10.1.111.111, 00:00:12
R       110.1.1.16 [120/1] via 10.1.111.111, 00:00:12
R       108.1.1.64 [120/1] via 10.1.111.111, 00:00:12
R       109.1.4.0 [120/1] via 10.1.111.111, 00:00:12
R       106.1.1.32 [120/1] via 10.1.111.111, 00:00:12
R       107.1.1.0 [120/1] via 10.1.111.111, 00:00:12
R       104.0.0.0 [120/1] via 10.1.111.111, 00:00:12
R       105.0.0.0 [120/1] via 10.1.111.111, 00:00:12
```

Note, even though the networks are in the routing table of R1, they are NOT reachable, because R1 does NOT have NLRI (Network Layer Reachability Information) to the next hop IP address (10.1.111.111) that is advertised. In order to provide reachability, the following static route is configured.

On R1

R1(config)#IP route 10.1.111.111 255.255.255.255 F0/0

To test and verify the configuration:

On R1

```
R1#Ping 101.0.0.111

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 101.0.0.111, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 110.1.1.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 110.1.1.17, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 7

Configure R1 such that only the existing and future prefixes with prefix-length of /10 to /26 are allowed in R1's routing table. R1 should receive these routes from BB1 and BB1 ONLY. Do not use neighbor command to accomplish this task.

On R1

The following prefix-list identifies the existing and the future routes that have a prefix-length of /10 through /26:

R1(config)#IP **prefix-list** NET **seq 5 permit** 0.0.0.0/0 **ge 10 le 26**

The following prefix-list identifies the BB1 router:

R1(config)#IP **prefix-list** BB1 **seq 5 permit 10.1.111.111/32**

R1(config)#Router rip
R1(config-router)#**Distribute-list prefix** NET **gateway** BB1 **in F0/0**

To test and verify the configuration:

On R1

R1#Show ip route | inc 10.1.111.111

```
R       103.0.0.0 [120/1] via 10.1.111.111, 00:00:25
S       10.1.111.111/32 is directly connected, FastEthernet0/0
R       108.1.1.64 [120/1] via 10.1.111.111, 00:00:25
R       109.1.4.0 [120/1] via 10.1.111.111, 00:00:25
R       107.1.1.0 [120/1] via 10.1.111.111, 00:00:25
R       104.0.0.0 [120/1] via 10.1.111.111, 00:00:25
R       105.0.0.0 [120/1] via 10.1.111.111, 00:00:25
```

Task 8

Configure Eigrp 100 on R2, and advertise it's link to BB2, if this configuration is done properly, R2 should receive two routes from BB2.

On R2

R2(config)#Router eigrp 100
R2(config-router)#No au
R2(config-router)#**Network 10.1.112.2 0.0.0.0**

To test and verify the configuration:

On R2

R2#Show ip route eigrp

```
     112.0.0.0/24 is subnetted, 2 subnets
D       112.2.2.0 [90/156160] via 10.1.112.112, 00:04:52, FastEthernet0/0
D       112.1.1.0 [90/156160] via 10.1.112.112, 00:04:52, FastEthernet0/0
```

Task 9

R2 should be configured to inject a default route into RIPv2's routing domain as long as any one of the two networks are in it's routing table.

On R2

R2(config)#**Access-list 1 permit 112.1.1.0 0.0.0.255**

R2(config)#**Access-list** 1 **permit 112.2.2.0 0.0.0.255**

R2(config)#**Route-map** TST **permit** 10
R2(config-route-map)#**Match ip addr** 1

R2(config)#**Router rip**
R2(config-router)#**Default-information originate route-map** TST

To verify the configuration:

On R6

R6#Show ip route rip | inc 0.0.0.0/0

R* 0.0.0.0/0 [120/2] via 10.1.46.4, 00:00:50, Serial0/0.64

Task 10

R1 should be configured such that R4 does not advertise the allowed networks from BB1 to its down stream neighbor/s.

On R1

R1(config)#**Access-list** 1 **permit 103.0.0.0 0.63.255.255**
R1(config)#**Access-list** 1 **permit 104.0.0.0 0.31.255.255**
R1(config)#**Access-list** 1 **permit 105.0.0.0 0.3.255.255**
R1(config)#**Access-list** 1 **permit 107.1.1.0 0.0.0.127**
R1(config)#**Access-list** 1 **permit 108.1.1.64 0.0.0.63**
R1(config)#**Access-list** 1 **permit 109.1.4.0 0.0.3.255**

R1(config)#Router rip
R1(config-router)#**Offset-list** 1 **out 12 Serial0/0.12**
R1(config-router)#**Offset-list** 1 **out 13 Serial0/0.14**

R2#Show ip route | B Gateway

```
Gateway of last resort is not set

     103.0.0.0/10 is subnetted, 1 subnets
R       103.0.0.0 [120/14] via 10.1.12.1, 00:00:44, Serial0/0.21
     100.0.0.0/24 is subnetted, 1 subnets
```

```
R       100.1.111.0 [120/1] via 10.1.12.1, 00:00:44, Serial0/0.21
     112.0.0.0/24 is subnetted, 2 subnets
D       112.2.2.0 [90/156160] via 10.1.112.112, 00:21:23, FastEthernet0/0
D       112.1.1.0 [90/156160] via 10.1.112.112, 00:21:23, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
R       10.1.14.0/24 [120/1] via 10.1.24.4, 00:00:02, FastEthernet0/1
                     [120/1] via 10.1.12.1, 00:00:44, Serial0/0.21
R       10.1.13.0/24 [120/1] via 10.1.12.1, 00:00:44, Serial0/0.21
C       10.1.12.0/24 is directly connected, Serial0/0.21
R       10.1.3.0/24 [120/2] via 10.1.12.1, 00:00:44, Serial0/0.21
R       10.1.111.111/32 [120/1] via 10.1.12.1, 00:00:44, Serial0/0.21
C       10.1.24.0/24 is directly connected, FastEthernet0/1
R       10.1.46.0/24 [120/1] via 10.1.24.4, 00:00:02, FastEthernet0/1
R       10.1.45.0/24 [120/1] via 10.1.24.4, 00:00:03, FastEthernet0/1
R       10.1.56.0/24 [120/2] via 10.1.24.4, 00:00:03, FastEthernet0/1
C       10.1.112.0/24 is directly connected, FastEthernet0/0
     108.0.0.0/26 is subnetted, 1 subnets
R       108.1.1.64 [120/14] via 10.1.12.1, 00:00:46, Serial0/0.21
     109.0.0.0/22 is subnetted, 1 subnets
R       109.1.4.0 [120/14] via 10.1.12.1, 00:00:47, Serial0/0.21
     107.0.0.0/25 is subnetted, 1 subnets
R       107.1.1.0 [120/14] via 10.1.12.1, 00:00:47, Serial0/0.21
     104.0.0.0/11 is subnetted, 1 subnets
R       104.0.0.0 [120/14] via 10.1.12.1, 00:00:47, Serial0/0.21
     105.0.0.0/14 is subnetted, 1 subnets
R       105.0.0.0 [120/14] via 10.1.12.1, 00:00:47, Serial0/0.21
```

Note R2 has a cost of 14 hops, these networks are advertised to R4 with a hop count of 15, which means that R4 can not advertise these networks to any other router.

On R4

R4#Show ip route | Inc 15

```
R       103.0.0.0 [120/15] via 10.1.24.2, 00:00:06, FastEthernet0/0
                  [120/15] via 10.1.14.1, 00:00:32, Serial0/0.41
R       108.1.1.64 [120/15] via 10.1.24.2, 00:00:06, FastEthernet0/0
                   [120/15] via 10.1.14.1, 00:00:32, Serial0/0.41
R       109.1.4.0 [120/15] via 10.1.24.2, 00:00:06, FastEthernet0/0
                  [120/15] via 10.1.14.1, 00:00:32, Serial0/0.41
R       107.1.1.0 [120/15] via 10.1.24.2, 00:00:06, FastEthernet0/0
                  [120/15] via 10.1.14.1, 00:00:32, Serial0/0.41
R       104.0.0.0 [120/15] via 10.1.24.2, 00:00:06, FastEthernet0/0
                  [120/15] via 10.1.14.1, 00:00:32, Serial0/0.41
```

```
R       105.0.0.0 [120/15] via 10.1.24.2, 00:00:06, FastEthernet0/0
                  [120/15] via 10.1.14.1, 00:00:32, Serial0/0.41
```

Task 11

R1 and R3 should be configured such that periodic RIPv2 updates are suppressed over the frame-relay connection between them. These routers should only send updates through the frame-relay connection if there is a topology change.

On R1

R1(config)#Int S0/0.13
R1(config-router)#**Ip rip triggered**

On R3

R3(config)#Int S0/0.31
R3(config-if)#**Ip rip triggered**

Note this command works on all point-to-point and some multipoint links.

Task 12

Configure R2 and R4 such that they exchange updates using Unicast.

On R2

R2(config)#Router rip
R2(config-router)#Passive-interface F0/1
R2(config-router)#Neighbor 10.1.24.4

On R4

R4(config)#Router rip
R4(config-router)#Passive-interface F0/0
R4(config-router)#Neighbor 10.1.24.2

Note if the "passive-interface" command is not used, the routers will send both Unicast and Multicast updates to each other.

Task 13

Configure R5 and R6 such that they exchange version 2 updates using Broadcast.

On R5 and R6

```
(config)#Int F0/0
(config-if)#Ip rip v2-broadcast
```

To test and verify the configuration:

On R6

R6#Debug ip rip

RIP: sending v2 flash update to **255.255.255.255** via FastEthernet0/0 (10.1.56.6)

Task 14

Configure the following Loopback interfaces on R3 and advertise a single summary route into the RIP routing domain:

Loopback 1 = 150.1.0.3 /24
Loopback 2 = 150.1.1.3 /24
Loopback 3 = 150.1.2.3 /24
Loopback 4 = 150.1.3.3 /24

On R3

```
R3(config)#Int Lo1
R3(config-if)#Ip address 150.1.0.3 255.255.255.0

R3(config)#Int lo2
R3(config-if)# Ip address 150.1.1.3 255.255.255.0
```

```
R3(config)#Int lo3
R3(config-if)#Ip address 150.1.2.3 255.255.255.0

R3(config)#Int lo4
R3(config-if)#Ip address 150.1.3.3 255.255.255.0

R3(config)#Router rip
R3(config-router)#Network 150.1.0.0

R3(config)#Int S0/0.31
R3(config-if)#Ip summary-address rip 150.1.0.0 255.255.252.0

R3(config)#Int S0/1
R3(config-if)#Ip summary-address rip 150.1.0.0 255.255.252.0
```

To test and verify the configuration:

On R1

R1#Show ip route rip

```
    103.0.0.0/10 is subnetted, 1 subnets
R       103.0.0.0 [120/1] via 10.1.111.111, 00:00:05
     10.0.0.0/8 is variably subnetted, 11 subnets, 2 masks
R       10.1.3.0/24    [120/1] via 10.1.100.3, 00:00:39, Serial0/1
                       [120/1] via 10.1.13.3, 00:00:56, Serial0/0.13
R       10.1.24.0/24  [120/1] via 10.1.14.4, 00:00:33, Serial0/0.14
                       [120/1] via 10.1.12.2, 00:00:49, Serial0/0.12
R       10.1.46.0/24   [120/1] via 10.1.14.4, 00:00:33, Serial0/0.14
R       10.1.45.0/24   [120/1] via 10.1.14.4, 00:00:33, Serial0/0.14
R       10.1.56.0/24   [120/2] via 10.1.14.4, 00:00:33, Serial0/0.14
R       10.1.112.0/24 [120/1] via 10.1.12.2, 00:00:49, Serial0/0.12
     108.0.0.0/26 is subnetted, 1 subnets
R       108.1.1.64 [120/1] via 10.1.111.111, 00:00:05
     109.0.0.0/22 is subnetted, 1 subnets
R       109.1.4.0 [120/1] via 10.1.111.111, 00:00:05
     107.0.0.0/25 is subnetted, 1 subnets
R       107.1.1.0 [120/1] via 10.1.111.111, 00:00:05
     104.0.0.0/11 is subnetted, 1 subnets
R       104.0.0.0 [120/1] via 10.1.111.111, 00:00:06
     150.1.0.0/22 is subnetted, 1 subnets
R       150.1.0.0 [120/1] via 10.1.100.3, 00:00:40, Serial0/1
                  [120/1] via 10.1.13.3, 00:00:57, Serial0/0.13
     105.0.0.0/14 is subnetted, 1 subnets
```

```
R       105.0.0.0 [120/1] via 10.1.111.111, 00:00:08
R*      0.0.0.0/0  [120/1] via 10.1.12.2, 00:00:51, Serial0/0.12
```

Task 15

R1 is a high speed router sending updates to R3 which is a low speed router. Because of this fact, R3 is not be able to receive and process updates at the rate that R1 operates. Configure R1 such that when it has multiple RIP packets to send to R3, it waits 10 milliseconds between the packets. To further remedy these situations, configure R3 to increase its unprocessed RIP input queue depth to 75 packets.

On R1

R1(config)#Router rip
R1(config-router)#**output-delay 10**

The above configuration will help R3 from losing routing information, because this command introduces a delay of 10 milliseconds between packets in a multiple packet RIP updates. By default there is no inter-packet delay and the range for this timer is (8 – 50 milliseconds).

On R3

R3(config)#Router rip
R3(config-router)#**input-queue 75**

This command will also help to prevent routing information from being lost. The value specifies the depth of the input queue, the larger the value, the larger the depth of the queue. The range is (0 – 1024) and the default value is 50.

Task 16

Configure R6 with the following 10 Loopback interfaces. R6 should be configured to advertise these Loopback interfaces in RIP routing domain.
Configure R6 such that R4 receives the EVEN routes from R6 and the ODD routes from R5. Whereas, R5 should receive the ODD routes from R6 and the EVEN numbered routes from R4. You should use an access-list with minimum number of lines to accomplish this task.

Loopback 0 = 160.1.0.6 /24, Loopback 1 = 160.1.1.6 /24, Loopback 2 = 160.1.2.6 /24
Loopback 3 = 160.1.3.6 /24, Loopback 4 = 160.1.4.6 /24, Loopback 5 = 160.1.5.6 /24
Loopback 6 = 160.1.6.6 /24, Loopback 7 = 160.1.7.6 /24, Loopback 8 = 160.1.8.6 /24
Loopback 9 = 160.1.9.6 /24.

On R6

```
R6(config)#Int lo0
R6(config-if)#Ip address 160.1.0.6 255.255.255.0

R6(config)#Int lo1
R6(config-if)# Ip address 160.1.1.6 255.255.255.0

R6(config)#Int lo2
R6(config-if)#Ip address 160.1.2.6 255.255.255.0

R6(config)#Int lo3
R6(config-if)#Ip address 160.1.3.6 255.255.255.0

R6(config)#Int lo4
R6(config-if)#Ip address 160.1.4.6 255.255.255.0

R6(config)#Int lo5
R6(config-if)#Ip address 160.1.5.6 255.255.255.0

R6(config)#Int lo6
R6(config-if)#Ip address 160.1.6.6 255.255.255.0

R6(config)#Int lo7
R6(config-if)#Ip address 160.1.7.6 255.255.255.0

R6(config)#Int lo8
R6(config-if)# Ip address 160.1.8.6 255.255.255.0

R6(config)#Int lo9
R6(config-if)#Ip address 160.1.9.6 255.255.255.0
```

On R6

```
R6(config)#Access-list 1 permit 160.1.1.0 0.0.254.255
R6(config)#Access-list 2 permit 160.1.0.0 0.0.254.255

R6(config)#Router rip
```

```
R6(config-router)#Offset-list 2 out 15 FastEthernet0/0
R6(config-router)#Offset-list 1 out 15 Serial0/0.64
R6(config-router)# Network 160.1.0.0
```

To verify and test the configuration:

On R4

```
R4#Sh ip route rip | Inc 10.1.46.6

R       160.1.0.0 [120/1] via 10.1.46.6, 00:00:29, Serial0/0.46
R       160.1.2.0 [120/1] via 10.1.46.6, 00:00:29, Serial0/0.46
R       160.1.4.0 [120/1] via 10.1.46.6, 00:00:29, Serial0/0.46
R       160.1.6.0 [120/1] via 10.1.46.6, 00:00:29, Serial0/0.46
R       160.1.8.0 [120/1] via 10.1.46.6, 00:00:29, Serial0/0.46
R       10.1.56.0/24 [120/1] via 10.1.46.6, 00:00:29, Serial0/0.46
```

```
R4#Sh ip route rip | Inc 10.1.45.5

R       160.1.1.0 [120/2] via 10.1.45.5, 00:00:51, Serial0/0.45
R       160.1.3.0 [120/2] via 10.1.45.5, 00:00:51, Serial0/0.45
R       160.1.5.0 [120/2] via 10.1.45.5, 00:00:51, Serial0/0.45
R       160.1.7.0 [120/2] via 10.1.45.5, 00:00:51, Serial0/0.45
R       160.1.9.0 [120/2] via 10.1.45.5, 00:00:51, Serial0/0.45
                  [120/1] via 10.1.45.5, 00:00:51, Serial0/0.45
```

Note R4 is receiving even subnets of 160.1.0.0 network from R6, whereas, the odd subnets of the same network is received from R5

On R5

```
R5#Show ip route rip | Inc 160.1.

     160.1.0.0/24 is subnetted, 10 subnets
R       160.1.1.0 [120/1] via 10.1.56.6, 00:00:26, FastEthernet0/0
R       160.1.0.0 [120/2] via 10.1.45.4, 00:00:26, Serial0/0.54
R       160.1.3.0 [120/1] via 10.1.56.6, 00:00:26, FastEthernet0/0
R       160.1.2.0 [120/2] via 10.1.45.4, 00:00:26, Serial0/0.54
R       160.1.5.0 [120/1] via 10.1.56.6, 00:00:26, FastEthernet0/0
R       160.1.4.0 [120/2] via 10.1.45.4, 00:00:26, Serial0/0.54
R       160.1.7.0 [120/1] via 10.1.56.6, 00:00:26, FastEthernet0/0
R       160.1.6.0 [120/2] via 10.1.45.4, 00:00:26, Serial0/0.54
R       160.1.9.0 [120/1] via 10.1.56.6, 00:00:26, FastEthernet0/0
R       160.1.8.0 [120/2] via 10.1.45.4, 00:00:26, Serial0/0.54
```

Note the even subnets of 160.1.0.0 are received from R4, whereas, the odd subnets of 160.1.0.0 are received from R6.

Task 17

Configure RIPv2 on BB3; this router is connected to CAT-1's port F0/13. Configure a solution such that R3 advertises all the RIPv2 routes to BB3. Do not change the VLAN assignment of any of the routers, or use a global configuration, and/or router configuration mode command to accomplish this task.
BB3 may not have reachability to any of the IP addresses within this topology.

On BB3

BB3(config)#Router rip
BB3(config-router)# No **validate-update-source**

On SW1

SW1(config)#Monitor session 1 source interface F0/3 both
SW1(config)#Monitor session 1 destination interface F0/13

To verify the configuration:

On BB3

BB3#Sh ip route rip

```
     103.0.0.0/10 is subnetted, 1 subnets

R       103.0.0.0 [120/3] via 10.1.3.3, 00:00:06
     100.0.0.0/24 is subnetted, 1 subnets
R       100.1.111.0 [120/2] via 10.1.3.3, 00:00:06
     160.1.0.0/24 is subnetted, 10 subnets
R       160.1.1.0 [120/5] via 10.1.3.3, 00:00:06
R       160.1.0.0 [120/4] via 10.1.3.3, 00:00:06
R       160.1.3.0 [120/5] via 10.1.3.3, 00:00:06
R       160.1.2.0 [120/4] via 10.1.3.3, 00:00:06
R       160.1.5.0 [120/5] via 10.1.3.3, 00:00:06
R       160.1.4.0 [120/4] via 10.1.3.3, 00:00:06
R       160.1.7.0 [120/5] via 10.1.3.3, 00:00:06
```

```
R       160.1.6.0 [120/4] via 10.1.3.3, 00:00:06
R       160.1.9.0 [120/5] via 10.1.3.3, 00:00:06
R       160.1.8.0 [120/4] via 10.1.3.3, 00:00:06
     10.0.0.0/8 is variably subnetted, 11 subnets, 2 masks
R       10.1.14.0/24 [120/2] via 10.1.3.3, 00:00:06
R       10.1.13.0/24 [120/1] via 10.1.3.3, 00:00:06
R       10.1.12.0/24 [120/2] via 10.1.3.3, 00:00:07
R       10.1.111.111/32 [120/2] via 10.1.3.3, 00:00:06
R       10.1.24.0/24 [120/3] via 10.1.3.3, 00:00:08
R       10.1.46.0/24 [120/3] via 10.1.3.3, 00:00:08
R       10.1.45.0/24 [120/3] via 10.1.3.3, 00:00:08
R       10.1.56.0/24 [120/4] via 10.1.3.3, 00:00:09
R       10.1.100.0/24 [120/1] via 10.1.3.3, 00:00:09
R       10.1.112.0/24 [120/3] via 10.1.3.3, 00:00:09
     108.0.0.0/26 is subnetted, 1 subnets
R       108.1.1.64 [120/3] via 10.1.3.3, 00:00:09
     109.0.0.0/22 is subnetted, 1 subnets
R       109.1.4.0 [120/3] via 10.1.3.3, 00:00:09
     107.0.0.0/25 is subnetted, 1 subnets
R       107.1.1.0 [120/3] via 10.1.3.3, 00:00:09
     104.0.0.0/11 is subnetted, 1 subnets
R       104.0.0.0 [120/3] via 10.1.3.3, 00:00:09
     150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
R       150.1.3.0/24 [120/1] via 10.1.3.3, 00:00:09
R       150.1.2.0/24 [120/1] via 10.1.3.3, 00:00:09
R       150.1.1.0/24 [120/1] via 10.1.3.3, 00:00:09
R       150.1.0.0/24 [120/1] via 10.1.3.3, 00:00:09
R       150.1.0.0/22 [120/3] via 10.1.3.3, 00:00:09
     105.0.0.0/14 is subnetted, 1 subnets
R       105.0.0.0 [120/3] via 10.1.3.3, 00:00:09
R*   0.0.0.0/0 [120/3] via 10.1.3.3, 00:00:09
```

Task 18

Erase the startup configuration and reload the routers before proceeding to the next protocol.

Advanced CCIE Routing & Switching 2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

EIGRP

Lab 1 – Eigrp Configuration

Lab Setup:

VLANs:

- F0/0 interface of BB1 and R1 should be configured in VLAN 11
- F0/0 interface of R3 should be configured in VLAN 3
- F0/0 interface of BB2 and R2 should be in VLAN 22
- F0/0 interface of R5 and R6 should be configured in VLAN 56
- F0/1 interface of R2 should be configured in VLAN 2
- F0/0 interface of R4 should be configured in VLAN 4

Frame-relay:

- R4 should be configured with two sub-interfaces in a point-to-point manner, one connecting R4 to R5 and the second one connecting R4 to R6.
- R5 and R6 should each be configured with a single point-to-point sub-interface connection to R4.
- R1 should be configured with three point-to-point sub-interfaces connecting it to routers R2, R3 and R4.
- Routers R2, R3 and R4 should be configured with a point-to-point frame-relay connection to R1.
- Configure the bandwidth of R4's point-to-point frame-relay connection to R5 to be 512 Kbps.

Trunking:

The trunking should be established between SW-1 and SW-2 using ports F0/19 and F0/20.

IP Addressing chart:

Router	Interface	Connecting to:	IP Address
R1	S0/0.12	R2	10.1.12.1 /24
	S0/0.13	R3	10.1.13.1 /24
	S0/0.14	R4	10.1.14.1 /24
	S0/1	R3	10.1.100.1 /24
	F0/0	BB1	10.1.111.1 /24
	Lo0	-	150.1.1.1 /24
	Lo1	-	1.1.0.1 /24
	Lo2	-	1.1.1.1 /24
	Lo3	-	1.1.2.1 /24
	Lo4	-	1.1.3.1 /24
R2	S0/0.21	R1	10.1.12.2 /24
	F0/0	BB2	10.1.112.2 /24
	F0/1	-	10.1.2.2 /24
	Lo0	-	150.1.2.2 /24
R3	S0/0.31	R1	10.1.13.3 /24
	S0/1	R1	10.1.100.3 /24
	F0/0	-	10.1.3.3 /24
	Lo0	-	150.1.3.3 /24
R4	S0/0.41	R1	10.1.14.4 /24
	S0/0.45	R5	10.1.45.4 /24
	S0/0.46	R6	10.1.46.4 /24
	F0/0	-	10.1.4.4 /24
	Lo0	-	150.1.4.4 /24
R5	S0/0.54	R4	10.1.45.5 /24
	F0/0	R6	10.1.56.5 /24
	Lo0	-	150.1.5.5 /24
R6	S0/0.56	R4	10.1.46.6 /24
	F0/0	R5	10.1.56.6 /24
	Lo0	-	150.1.6.6 /24
BB1	F0/0	R1	10.1.111.111 /24
	Lo0	-	150.1.111.111 /24
BB2	F0/0	R2	10.1.112.112 /24
	Lo0	-	150.1.112.112 /24

Logical Topology

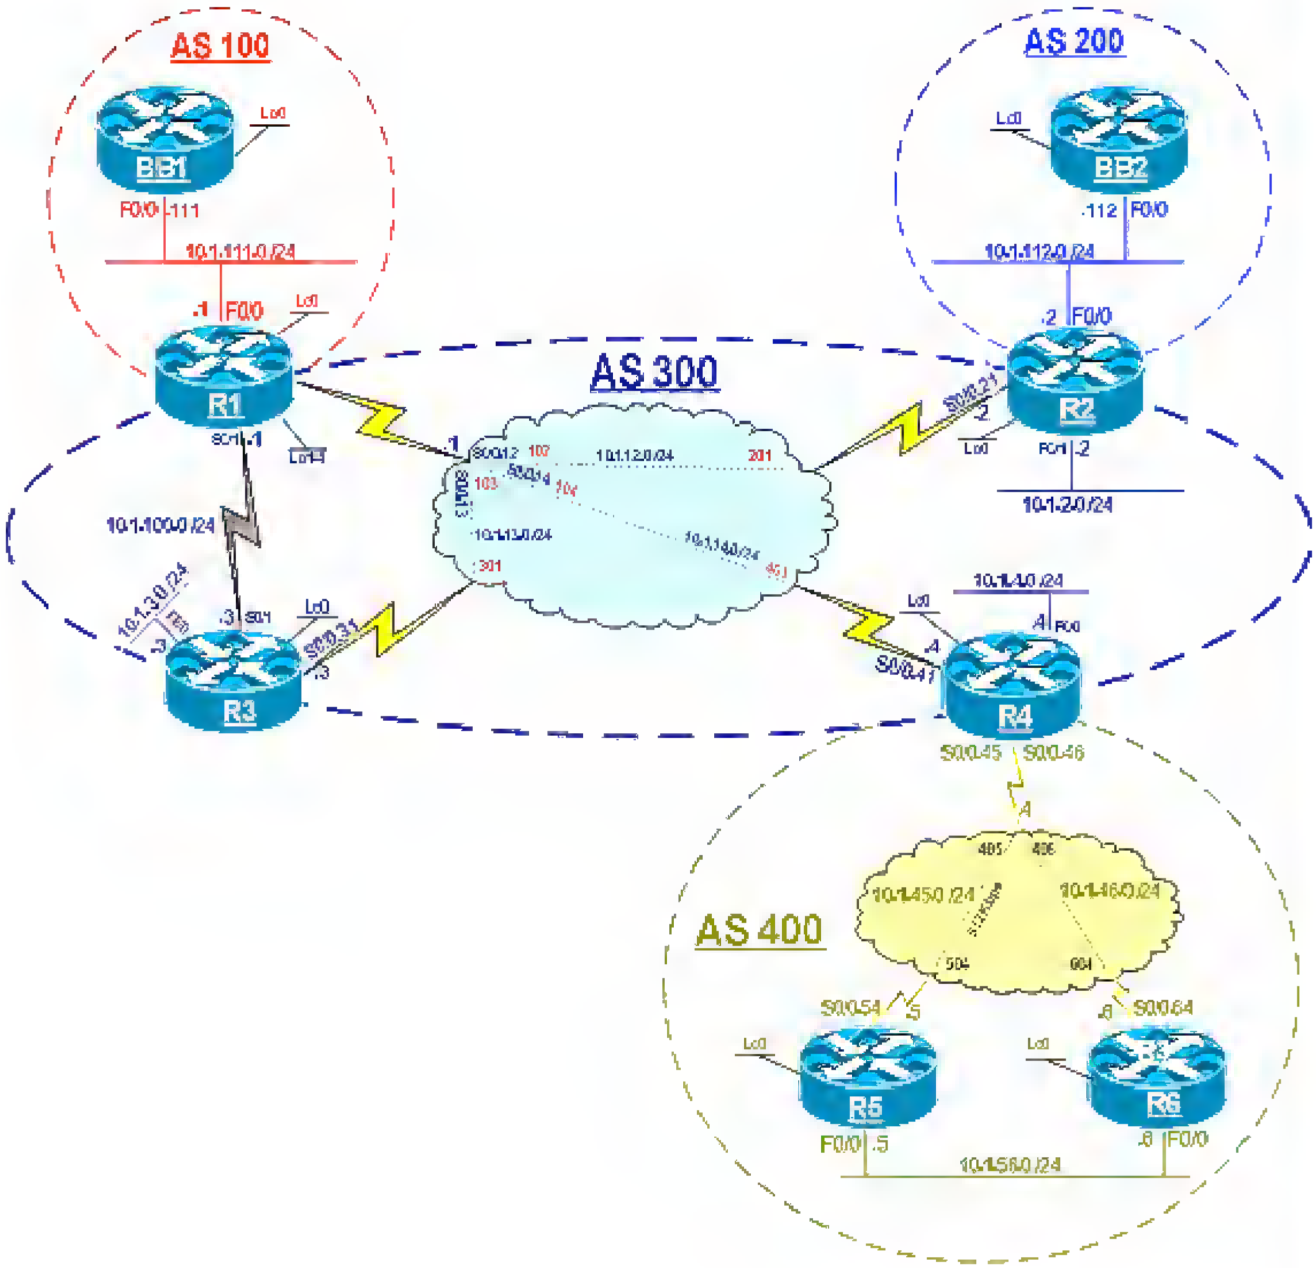

Task 1

Configure the routers as follows:

- R1 and BB1 should be configured in AS 100; BB1 should advertise it's directly connected networks in this AS, whereas, R1 should ONLY advertise it's connection to BB1 and it's lo0 interface in this AS.

- R2 and BB2 should be configured in AS 200; BB2 should advertise it's directly connected networks in this AS, whereas, R2 should ONLY advertise it's connection to BB2 in this AS.

- R1, R2, R3 and R4 should be configured in AS 300; R1 should advertise it's P2P connection to R3, all of it's frame-relay connections and Loopback 1 – 4 in this AS. R2 should advertise it's F0/1, Lo0 and it's frame-relay connection in this AS. R3 should advertise all of it's interfaces in this AS. R4 should advertise it's frame-relay connection to R1, F0/0 and it's Lo0 interface in this AS.

- R4, R5 and R6 should be configured in AS 400; R4 should advertise it's frame-relay connections to R5 and R6 in this AS. R5 and R6 should advertise all their directly connected networks in this AS.

To configure the first item in this task:

On BB1

```
BB1(config)#router eigrp 100
BB1(config-router)#no au
BB1(config-router)#network 0.0.0.0
```

Note the "Network 0.0.0.0" advertises the existing and future directly connected networks in the AS.

On R1

```
R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#network 10.1.111.1 0.0.0.0
R1(config-router)#network 150.1.1.1 0.0.0.0
```

To verify the configuration:

On BB1

BB1#Show ip route eigrp

```
     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.1.0 [90/156160] via 10.1.111.1, 00:02:06, FastEthernet0/0
```

On R1

R1#Show ip route eigrp

```
     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.111.0 [90/156160] via 10.1.111.111, 00:02:15, FastEthernet0/0
```

To configure the second item in this task:

On BB2

```
BB2(config)#router eigrp 200
BB2(config-router)#no au
BB2(config-router)#network 0.0.0.0
```

On R2

```
R2(config)#router eigrp 200
R2(config-router)#no au
R2(config-router)#network 10.1.112.0 0.0.0.255
```

Note the above network command is another way to advertise routes in Eigrp.

To verify the configuration:

On BB2

BB2#Show ip eigrp neighbors

```
IP-EIGRP neighbors for process 200
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q   Seq
                                            (sec)         (ms)       Cnt  Num
0   10.1.112.2              Fa0/0             11 00:00:30    4   200  0   2
```

BB2#Show ip eigrp topology

```
IP-EIGRP Topology Table for AS(200)/ID(150.1.112.112)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
```

```
       r - reply Status, s - sia Status

P 150.1.112.0/24, 1 successors, FD is 128256
         via Connected, Loopback0

P 10.1.112.0/24, 1 successors, FD is 28160
         via Connected, FastEthernet0/0
```

On R2

R2#Show ip route eigrp

```
     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.112.0 [90/156160] via 10.1.112.112, 00:03:47, FastEthernet0/0
```

R2#Show ip eigrp topology

```
IP-EIGRP Topology Table for AS(200)/ID(150.1.2.2)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 150.1.112.0/24, 1 successors, FD is 156160
        via 10.1.112.112 (156160/128256), FastEthernet0/0
P 10.1.112.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
```

R2#Show ip eigrp topology 150.1.112.0/24

```
IP-EIGRP (AS 200): Topology entry for 150.1.112.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 156160
  Routing Descriptor Blocks:
  10.1.112.112 (FastEthernet0/0), from 10.1.112.112, Send flag is 0x0
      Composite metric is (156160/128256), Route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
```

To configure the third item in this task:

On R1

R1(config)#router eigrp 300
R1(config-router)#no au
R1(config-router)#network 10.1.100.1 0.0.0.0
R1(config-router)#network 10.1.12.1 0.0.0.0
R1(config-router)#network 10.1.13.1 0.0.0.0
R1(config-router)#network 10.1.14.1 0.0.0.0
R1(config-router)#network 1.1.0.1 0.0.0.0
R1(config-router)#network 1.1.1.1 0.0.0.0
R1(config-router)#network 1.1.2.1 0.0.0.0
R1(config-router)#network 1.1.3.1 0.0.0.0

On R2

R2(config)#router eigrp 300
R2(config-router)#no au
R2(config-router)#network 10.1.2.2 0.0.0.0
R2(config-router)#network 150.1.2.2 0.0.0.0
R2(config-router)#network 10.1.12.2 0.0.0.0

On R3

R3(config)#router eigrp 300
R3(config-router)#no au
R3(config-router)#network 0.0.0.0

On R4

R4(config)#router eigrp 300
R4(config-router)#no au
R4(config-router)#network 10.1.14.4 0.0.0.0
R4(config-router)#network 150.1.4.4 0.0.0.0
R4(config-router)#network 10.1.4.4 0.0.0.0

To verify the configuration:

On R1

R1#Sh ip route eigrp 300

     10.0.0.0/24 is subnetted, 8 subnets
D   10.1.3.0 [90/2172416] via 10.1.100.3, 00:02:02, Serial0/1

```
                    [90/2172416] via 10.1.13.3, 00:02:02, Serial0/0.13
D       10.1.2.0 [90/2172416] via 10.1.12.2, 00:06:29, Serial0/0.12
D       10.1.4.0 [90/2172416] via 10.1.14.4, 00:04:33, Serial0/0.14
    150.1.0.0/24 is subnetted, 5 subnets
D       150.1.4.0 [90/2297856] via 10.1.14.4, 00:04:23, Serial0/0.14
D       150.1.3.0 [90/2297856] via 10.1.100.3, 00:02:02, Serial0/1
                    [90/2297856] via 10.1.13.3, 00:02:02, Serial0/0.13
D       150.1.2.0 [90/2297856] via 10.1.12.2, 00:06:29, Serial0/0.12
```

On R2

R2#Show ip route eigrp 300

```
        1.0.0.0/24 is subnetted, 4 subnets
D       1.1.0.0 [90/2297856] via 10.1.12.1, 00:07:21, Serial0/0.21
D       1.1.1.0 [90/2297856] via 10.1.12.1, 00:07:21, Serial0/0.21
D       1.1.2.0 [90/2297856] via 10.1.12.1, 00:07:21, Serial0/0.21
D       1.1.3.0 [90/2297856] via 10.1.12.1, 00:07:21, Serial0/0.21
    10.0.0.0/24 is subnetted, 8 subnets
D       10.1.14.0 [90/2681856] via 10.1.12.1, 00:07:21, Serial0/0.21
D       10.1.13.0 [90/2681856] via 10.1.12.1, 00:07:21, Serial0/0.21
D       10.1.3.0 [90/2684416] via 10.1.12.1, 00:06:03, Serial0/0.21
D       10.1.4.0 [90/2684416] via 10.1.12.1, 00:05:26, Serial0/0.21
D       10.1.100.0 [90/2681856] via 10.1.12.1, 00:07:21, Serial0/0.21
    150.1.0.0/24 is subnetted, 4 subnets
D       150.1.4.0 [90/2809856] via 10.1.12.1, 00:05:16, Serial0/0.21
D       150.1.3.0 [90/2809856] via 10.1.12.1, 00:06:03, Serial0/0.21
```

On R3

R3#Sh ip route eigrp 300

```
        1.0.0.0/24 is subnetted, 4 subnets
D       1.1.0.0 [90/2297856] via 10.1.100.1, 00:04:53, Serial0/1
                [90/2297856] via 10.1.13.1, 00:04:53, Serial0/0.31
D       1.1.1.0 [90/2297856] via 10.1.100.1, 00:04:53, Serial0/1
                [90/2297856] via 10.1.13.1, 00:04:53, Serial0/0.31
D       1.1.2.0 [90/2297856] via 10.1.100.1, 00:04:53, Serial0/1
                [90/2297856] via 10.1.13.1, 00:04:53, Serial0/0.31
D       1.1.3.0 [90/2297856] via 10.1.100.1, 00:04:53, Serial0/1
                [90/2297856] via 10.1.13.1, 00:04:53, Serial0/0.31
    10.0.0.0/24 is subnetted, 7 subnets
D       10.1.14.0 [90/2681856] via 10.1.100.1, 00:04:53, Serial0/1
                  [90/2681856] via 10.1.13.1, 00:04:53, Serial0/0.31
```

```
D       10.1.12.0 [90/2681856] via 10.1.100.1, 00:04:53, Serial0/1
                  [90/2681856] via 10.1.13.1, 00:04:53, Serial0/0.31
D       10.1.2.0 [90/2684416] via 10.1.100.1, 00:04:53, Serial0/1
                 [90/2684416] via 10.1.13.1, 00:04:53, Serial0/0.31
D       10.1.4.0 [90/2684416] via 10.1.100.1, 00:04:54, Serial0/1
                 [90/2684416] via 10.1.13.1, 00:04:54, Serial0/0.31
     150.1.0.0/24 is subnetted, 3 subnets
D       150.1.4.0 [90/2809856] via 10.1.100.1, 00:04:54, Serial0/1
                  [90/2809856] via 10.1.13.1, 00:04:54, Serial0/0.31
D       150.1.2.0 [90/2809856] via 10.1.100.1, 00:04:54, Serial0/1
                  [90/2809856] via 10.1.13.1, 00:04:54, Serial0/0.31
```

On R4

R4#Show ip route eigrp 300

```
   1.0.0.0/24 is subnetted, 4 subnets
D       1.1.0.0 [90/2297856] via 10.1.14.1, 00:09:11, Serial0/0.41
D       1.1.1.0 [90/2297856] via 10.1.14.1, 00:09:11, Serial0/0.41
D       1.1.2.0 [90/2297856] via 10.1.14.1, 00:09:11, Serial0/0.41
D       1.1.3.0 [90/2297856] via 10.1.14.1, 00:09:11, Serial0/0.41
     10.0.0.0/24 is subnetted, 9 subnets
D       10.1.13.0 [90/2681856] via 10.1.14.1, 00:09:11, Serial0/0.41
D       10.1.12.0 [90/2681856] via 10.1.14.1, 00:09:11, Serial0/0.41
D       10.1.3.0 [90/2684416] via 10.1.14.1, 00:09:11, Serial0/0.41
D       10.1.2.0 [90/2684416] via 10.1.14.1, 00:09:11, Serial0/0.41
D       10.1.100.0 [90/2681856] via 10.1.14.1, 00:09:11, Serial0/0.41
     150.1.0.0/24 is subnetted, 3 subnets
D       150.1.3.0 [90/2809856] via 10.1.14.1, 00:09:11, Serial0/0.41
D       150.1.2.0 [90/2809856] via 10.1.14.1, 00:09:11, Serial0/0.41
```

To configure the forth item in this task:

On R4

```
R4(config)#router eigrp 400
R4(config-router)#no au
R4(config-router)#network 10.1.45.4 0.0.0.0
R4(config-router)#network 10.1.46.4 0.0.0.0
```

On R5

```
R5(config)#router eigrp 400
R5(config-router)#no au
```

R5(config-router)#network 0.0.0.0

<u>On R6</u>

R6(config)#router eigrp 400
R6(config-router)#no au
R6(config-router)#network 0.0.0.0

<u>To verify the configuration:</u>

<u>On R4</u>

<u>R4#Show ip route eigrp 400</u>

```
     10.0.0.0/24 is subnetted, 10 subnets
D       10.1.56.0 [90/2172416] via 10.1.46.6, 00:00:32, Serial0/0.46
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.6.0 [90/2297856] via 10.1.46.6, 00:00:32, Serial0/0.46
D       150.1.5.0 [90/2300416] via 10.1.46.6, 00:00:32, Serial0/0.46
```

<u>On R5</u>

<u>R5#Show ip route eigrp</u>

```
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.46.0 [90/2172416] via 10.1.56.6, 00:01:03, FastEthernet0/0
     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.6.0 [90/156160] via 10.1.56.6, 00:01:00, FastEthernet0/0
```

<u>On R6</u>

<u>R6#Show ip route eigrp</u>

```
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.45.0 [90/2172416] via 10.1.56.5, 00:01:22, FastEthernet0/0
     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.5.0 [90/156160] via 10.1.56.5, 00:01:22, FastEthernet0/0
```

Task 2

Configure the hello and dead interval of all the routers in AS 300 to 20 and 80 respectively.

On R1

R1(config)#**Int S0/0.12**
R1(config-subif)#**ip hello-interval eigrp 300 20**
R1(config-subif)#**ip hold-time eigrp 300 80**

R1(config-subif)#**Int S0/0.13**
R1(config-subif)#**ip hello-interval eigrp 300 20**
R1(config-subif)#**ip hold-time eigrp 300 80**

R1(config-subif)#**Int S0/0.14**
R1(config-subif)#**ip hello-interval eigrp 300 20**
R1(config-subif)#**ip hold-time eigrp 300 80**

R1(config-subif)#**Int S0/1**
R1(config-if)#**ip hello-interval eigrp 300 20**
R1(config-if)#**ip hold-time eigrp 300 80**

On R2

R2(config)#**Int S0/0.21**
R2(config-subif)#**ip hello-interval eigrp 300 20**
R2(config-subif)#**ip hold-time eigrp 300 80**

On R3

R3(config)#**Int S0/0.31**
R3(config-subif)#**ip hello-interval eigrp 300 20**
R3(config-subif)#**ip hold-time eigrp 300 80**

R3(config-subif)#**Int S0/1**
R3(config-if)#**ip hello-interval eigrp 300 20**
R3(config-if)#**ip hold-time eigrp 300 80**

On R4

R4(config)#**Int S0/0.41**
R4(config-subif)#**ip hello-interval eigrp 300 20**
R4(config-subif)#**ip hold-time eigrp 300 80**

Task 3

Ensure that the routers in AS100 ONLY use bandwidth to calculate their composite metric.

Note the composite metric for network 150.1.111.0 /24 is calculated as follows:

10,000,000 Kbit divided by the slowest bandwidth along the path to a given destination (In this case network 150.1.111.0 /24), plus the sum of all interface delays along the path to that destination divided by 10 , and then, the result of the previous calculation should be multiplied by 256:

BB1#Show int lo0 | Inc MTU

```
  MTU 1514 bytes, BW 8000000 Kbit/sec, DLY 5000 usec,
```

BB1#Sh int F0/0

```
MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec
```

10,000,000 / 100,000 = 100

5000 + 100 = 5100/10 = 510

(100 + 510) * 256 = 156160   This is the composite metric that should be seen in the routing table

On R1

R1#sh ip route 150.1.111.0 | Inc metric

```
  Known via "eigrp 100", distance 90, metric 156160, type internal
      Route metric is 156160, traffic share count is 1
```

To change the K values based on the requirement:

On Both routers:

```
(config)#router eigrp 100
(config-router)#metric weight 0 1 0 0 0 0
```

Note once the K value of a router is changed, the neighbor adjacency goes down, because if the K values are different between two routers, the routers will not form neighbor adjacency. The following parameters must be the same

on two routers before they become adjacent:

- The K values
- AS numbers
- They must share the same layer two data link and be from the same IP address space.
- If authentication must be enabled, it must be enabled on both routers and the password for the authentication must match.

To verify the configuration:

On R1

R1#Show ip protocols

```
Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=0, K4=0, K5=0
```

(The rest of the output is omitted)

R1#Show ip eigrp 100 neighbors

```
IP-EIGRP neighbors for process 100
H   Address          Interface    Hold Uptime   SRTT   RTO  Q    Seq
                                  (sec)         (ms)        Cnt  Num
0   10.1.111.111     Fa0/0         11 00:10:34   5     300   0   8
```

R1#Sh ip rout eigrp 100

```
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.111.0 [90/25600] via 10.1.111.111, 00:12:22, FastEthernet0/0
```

Note once the routers are configured, the composite value is changed based on bandwidth multiplied by 256 as follows:

100 * 256 = 25600

Task 4

Ensure that the routers in AS200 ONLY use the delay parameter to calculate their composite metric.

On Both routers:

```
(config)#router eigrp 200
(config-router)#metric weight 0 0 0 1 0 0
```

Note in this case only the delay value is considered, therefore, the sum of all the interface delays divided by 10 should be multiplied by 256 as follows:

(5000 + 100) / 10 = 510 This is the sum of all interface delay values divided by 10.

510 * 256 = 130560 This should be the new composite value.

To test the configuration:

On R2

```
R2#Show ip route eigrp 200

     150.1.0.0/24 is subnetted, 4 subnets
D       150.1.112.0 [90/130560] via 10.1.112.112, 00:01:00, FastEthernet0/0
```

Task 5

Configure R1 to summarize it's Loopback 1 – 4 based on the following policy:

- R1 should ONLY advertise the summary route to R2.
- R1 should advertise the summary route plus the network for Loopback 2 to R3.
- R1 should advertise the summary route plus all the specific networks to R4.
- Only one summary command per neighbor should be used to accomplish this task.

To configure the first item:

On R1

R1(config)#int S0/0.12
R1(config-subif)#ip summary-address eigrp 300 1.1.0.0 255.255.252.0

You should see the following message:

%DUAL-5-NBRCHANGE: IP-EIGRP(0) 300: Neighbor 10.1.12.2 (Serial0/0.12) is resync: summary configured

To test the configuration:

On R2

R2#Show ip route eigrp 300 | Inc /22

1.0.0.0/22 is subnetted, 1 subnets
D 1.1.0.0 [90/2297856] via 10.1.12.1, 00:01:03, Serial0/0.21

To configure the second item:

Since configuring multiple summary commands is not allowed, one way to accomplish this task is to configure two virtual template interfaces, one between R1 and R3, and another one between R1 and R4.
Note the "leak-map" option is available under the physical and virtual-template interfaces.

On R1

To configure PPP, a virtual-template interface must be configured and the IP address of the sub-interface must be assigned to the virtual-template:

R1(config)#int S0/0.13
R1(config-subif)#no ip addr

R1(config)#**Int virtual-template 13**
R1(config-if)#ip address 10.1.13.1 255.255.255.0
R1(config-if)#**ip summary-address eigrp 300** 1.1.0.0 **255.255.252.0 leak-map R1-3**

R1(config)#**Route-map R1-3 permit** 10
R1(config-route-map)#**match ip addr 1**

R1(config)#**access-list 1 permit** 1.1.1.0 **0.0.0.255**

Note the leak-map option is now available, this option references a route-map, and the route-map references an access-list and what ever network/s that is permitted in the access-list will be leaked along the summary route.

Lastly the virtual-template interface is assigned to the sub-interface.

```
R1(config)#int S0/0.13
R1(config-subif)#frame-relay interface-dlci 103 ppp virtual-Template 13
```

Since there are two links (Int S0/0.13 and S0/1) between R1 and R3, the summary should also be applied to S0/1 interface:

```
R1(config)#Int S0/1
R1(config-if)#ip summary-address eigrp 300 1.1.0.0 255.255.252.0 leak-map R1-3
```

On R3

```
R3(config)#int S0/0.31
R3(config-subif)#no ip addr

R3(config)#Int virtual-template 31
R3(config-if)#ip address 10.1.13.3 255.255.255.0

R3(config-if)#int S0/0.31
R3(config-subif)#frame-relay interface-dlci 301 ppp virtual-Template 31
```

To verify the configuration:

On R3

```
R3#Sh ip route eigrp 300

     1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       1.1.0.0/22 [90/2297856] via 10.1.100.1, 00:02:15, Serial0/1
D       1.1.1.0/24 [90/2297856] via 10.1.100.1, 00:02:15, Serial0/1
     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
D       10.1.14.0/24 [90/2681856] via 10.1.100.1, 00:02:15, Serial0/1
D       10.1.12.0/24 [90/2681856] via 10.1.100.1, 00:02:15, Serial0/1
D       10.1.2.0/24 [90/2684416] via 10.1.100.1, 00:02:15, Serial0/1
D       10.1.4.0/24 [90/2684416] via 10.1.100.1, 00:02:15, Serial0/1
     150.1.0.0/24 is subnetted, 3 subnets
D       150.1.4.0 [90/2809856] via 10.1.100.1, 00:02:15, Serial0/1
D       150.1.2.0 [90/2809856] via 10.1.100.1, 00:02:15, Serial0/1
```

To configure the third item:

On R1

```
R1(config-subif)#int S0/0.14
R1(config-subif)#no ip addr

R1(config)#Int virtual-template 14
R1(config-if)#ip address 10.1.14.1 255.255.255.0
R1(config-if)#ip summary-address eigrp 300 1.1.0.0 255.255.252.0 leak-map R1-4

R1(config)#int S0/0.14
R1(config-subif)#frame-relay interface-dlci 104 ppp virtual-Template 14

R1(config)#Route-map R1-4 permit 10
```

Note if the leak-map references a route-map, and the route-map does not reference an access-list or it references an access-list that does not exist, the summary plus all specific routes are advertised.

On R4

```
R4(config)#int S0/0.41
R4(config-subif)#no ip addr

R4(config)#Int virtual-template 41
R4(config-if)#ip address 10.1.14.4 255.255.255.0

R4(config-if)#int S0/0.41
R4(config-subif)#frame-relay interface-dlci 401 ppp virtual-Template 41
```

To verify the configuration:

On R4

R4#Show ip route eigrp 300

```
     1.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D       1.1.0.0/24 [90/2713600] via 10.1.14.1, 00:00:13, Virtual-Access2
D       1.1.0.0/22 [90/2713600] via 10.1.14.1, 00:00:13, Virtual-Access2
D       1.1.1.0/24 [90/2713600] via 10.1.14.1, 00:00:13, Virtual-Access2
D       1.1.2.0/24 [90/2713600] via 10.1.14.1, 00:00:13, Virtual-Access2
D       1.1.3.0/24 [90/2713600] via 10.1.14.1, 00:00:13, Virtual-Access2
     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
```

```
D       10.1.13.3/32 [90/5145600] via 10.1.14.1, 00:00:13, Virtual-Access2
D       10.1.13.0/24 [90/5145600] via 10.1.14.1, 00:00:13, Virtual-Access2
D       10.1.12.0/24 [90/4729856] via 10.1.14.1, 00:00:13, Virtual-Access2
D       10.1.3.0/24 [90/4732416] via 10.1.14.1, 00:00:13, Virtual-Access2
D       10.1.2.0/24 [90/4732416] via 10.1.14.1, 00:00:13, Virtual-Access2
D       10.1.100.0/24 [90/4729856] via 10.1.14.1, 00:00:13, Virtual-Access2
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.3.0 [90/4857856] via 10.1.14.1, 00:00:14, Virtual-Access2
D       150.1.2.0 [90/4857856] via 10.1.14.1, 00:00:14, Virtual-Access2
```

Task 6

R4 should perform unequal cost load balancing to get to network 10.1.56.0 /24.

Note R4 takes R6 (10.1.46.6) to get to network 10.1.56.0 /24, the routing table of R4 reveals this information:

R4#Show ip route eigrp 400

```
     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D       10.1.56.0/24 [90/2172416] via 10.1.46.6, 00:44:40, Serial0/0.46
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.6.0 [90/2297856] via 10.1.46.6, 00:44:40, Serial0/0.46
D       150.1.5.0 [90/2300416] via 10.1.46.6, 00:44:40, Serial0/0.46
```

In order to perform an unequal cost load balancing, the advertised distance of the worst route should be lower than the feasible distance. In this case the advertise distance of R5 for network 10.1.56.0 /24 is 28160, this value is less than the feasible distance which is 2172416, this means that R5 meets the feasibility condition, therefore, the unequal cost load balancing can be performed.

R4#Show ip eigrp 400 topology 10.1.56.0/24

```
IP-EIGRP (AS 400): Topology entry for 10.1.56.0/24
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2172416
  Routing Descriptor Blocks:
  10.1.46.6 (Serial0/0.46), from 10.1.46.6, Send flag is 0x0
      Composite metric is (2172416/28160), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20100 microseconds
```

```
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
  10.1.45.5 (Serial0/0.45), from 10.1.45.5, Send flag is 0x0
    Composite metric is (5514496/28160), Route is Internal
    Vector metric:
      Minimum bandwidth is 512 Kbit
      Total delay is 20100 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
```

The last step in accomplishing this task is to divide the worst route by the best route to get the ratio:

5,514,496 / 2,172,416 = 2.538

The result is the number that must be configured using the "variance" command. This value should be rounded up; in this case the result is 3.

On R4

```
R4(config)#router eigrp 400
R4(config-router)#variance 3
```

To verify the configuration:

On R4

```
R4#Show ip route eigrp 400

     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D       10.1.56.0/24 [90/2172416] via 10.1.46.6, 00:00:23, Serial0/0.46
                     [90/5514496] via 10.1.45.5, 00:00:23, Serial0/0.45
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.6.0 [90/2297856] via 10.1.46.6, 00:00:23, Serial0/0.46
                  [90/5642496] via 10.1.45.5, 00:00:23, Serial0/0.45
D       150.1.5.0 [90/2300416] via 10.1.46.6, 00:00:23, Serial0/0.46
                  [90/5639936] via 10.1.45.5, 00:00:23, Serial0/0.45
```

Task 7

Configure R1 to disable the SIA timer for AS 300 and set the SIA timer to 60 minutes for AS 100.

On R1

R1(config)#**Router eigrp** 100
R1(config-router)#**timers active-time** 60

R1(config-router)#**Router eigrp** 300
R1(config-router)#**timers active-time disabled**

Task 8

Configure authentication for all the routers in AS 300 and set the passwords as follows:

- R1 and R2 should use "**Cisco12**".
- R1 and R3 should use "**Cisco13**".
- R1 and R4 should use "**Cisco14**".

To configure authentication between R1 and R2:

On R1 and R2

(config)#Key chain **R1-2**
(config-keychain)# key 1
(config-keychain -key)#key-string **Cisco12**

On R1

R1(config)#int S0/0.12
R1(config-subif)#ip authentication key-chain eigrp 300 **R1-2**
R1(config-subif)#ip authentication mode eigrp 300 md5

On R2

R2(config)#int S0/0.21

```
R2(config-subif)#ip authentication key-chain eigrp 300 R1-2
R2(config-subif)#ip authentication mode eigrp 300 md5
```

To test the configuration:

On R2

R2#Show ip route eigrp 300

```
     1.0.0.0/22 is subnetted, 1 subnets
D       1.1.0.0 [90/2297856] via 10.1.12.1, 00:00:19, Serial0/0.21
     10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D       10.1.14.4/32 [90/4729856] via 10.1.12.1, 00:00:19, Serial0/0.21
D       10.1.14.0/24 [90/4729856] via 10.1.12.1, 00:00:19, Serial0/0.21
D       10.1.13.3/32 [90/4729856] via 10.1.12.1, 00:00:19, Serial0/0.21
D       10.1.13.0/24 [90/4729856] via 10.1.12.1, 00:00:19, Serial0/0.21
D       10.1.3.0/24 [90/2684416] via 10.1.12.1, 00:00:19, Serial0/0.21
D       10.1.4.0/24 [90/4732416] via 10.1.12.1, 00:00:19, Serial0/0.21
D       10.1.100.0/24 [90/2681856] via 10.1.12.1, 00:00:19, Serial0/0.21
     150.1.0.0/24 is subnetted, 4 subnets
D       150.1.4.0 [90/4857856] via 10.1.12.1, 00:00:19, Serial0/0.21
D       150.1.3.0 [90/2809856] via 10.1.12.1, 00:00:19, Serial0/0.21
```

A "Show ip eigrp interface detail" command can also be used to verify the authentication.

R2#Show ip eigrp inter detail | B Se0/0.21

```
Se0/0.21          1       0/0       109       0/15        575          0
  Hello interval is 20 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 39/33
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 14
  Retransmissions sent: 3  Out-of-sequence rcvd: 3
  Authentication mode is md5,  key-chain is "R1-2"
  Use unicast
```

To configure authentication between R1 and R3:

On R1 and R3:

```
(config)#key chain R1-3
(config-keychain)#key 1
(config-keychain-key)#key-string Cisco13
```

On R1

```
R1(config)#int virtual-template13
R1(config-if)#ip authentication key-chain eigrp 300 R1-3
R1(config-if)#ip authentication mode eigrp 300 md5

R1(config-if)#int S0/1
R1(config-if)#ip authentication key-chain eigrp 300 R1-3
R1(config-if)#ip authentication mode eigrp 300 md5
```

On R3

```
R3(config)#int virtual-template31
R3(config-if)#ip authentication key-chain eigrp 300 R1-3
R3(config-if)#ip authentication mode eigrp 300 md5

R3(config-keychain-key)#int S0/1
R3(config-if)#ip authentication key-chain eigrp 300 R1-3
R3(config-if)#ip authentication mode eigrp 300 md5
```

To test the configuration:

On R3

R3#Show ip route eigrp 300

```
     1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       1.1.0.0/22 [90/2297856] via 10.1.100.1, 00:00:17, Serial0/1
D       1.1.1.0/24 [90/2297856] via 10.1.100.1, 00:00:17, Serial0/1
     10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D       10.1.14.4/32 [90/4729856] via 10.1.100.1, 00:00:17, Serial0/1
D       10.1.14.0/24 [90/4729856] via 10.1.100.1, 00:00:17, Serial0/1
D       10.1.12.0/24 [90/2681856] via 10.1.100.1, 00:00:17, Serial0/1
D       10.1.2.0/24 [90/2684416] via 10.1.100.1, 00:00:17, Serial0/1
D       10.1.4.0/24 [90/4732416] via 10.1.100.1, 00:00:17, Serial0/1
     150.1.0.0/24 is subnetted, 3 subnets
D       150.1.4.0 [90/4857856] via 10.1.100.1, 00:00:17, Serial0/1
D       150.1.2.0 [90/2809856] via 10.1.100.1, 00:00:17, Serial0/1
```

R3#Show ip eigrp 300 neighbors

```
IP-EIGRP neighbors for process 300

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q   Seq
```

```
                                    (sec)          (ms)          Cnt  Num
1   10.1.13.1         Vi2           10 00:00:27  132       792   0   227
0   10.1.100.1        Se0/1         66 00:06:08   32       200   0   223
```

To configure authentication between R1 and R4:

On R1 and R4:

```
(config)#key chain R1-4
(config-keychain)#key 1
(config-keychain-key)#key-string Cisco14
```

On R1

```
R1(config)#int virtual-template14
R1(config-if)#ip authentication key-chain eigrp 300 R1-4
R1(config-if)#ip authentication mode eigrp 300 md5
```

On R4

```
R4(config)#int virtual-template41
R4(config-if)#ip authentication key-chain eigrp 300 R1-4
R4(config-if)#ip authentication mode eigrp 300 md5
```

To verify the configuration:

On R4

R4#Show ip route eigrp 300

```
     1.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D       1.1.0.0/24 [90/2713600] via 10.1.14.1, 00:00:27, Virtual-Access2
D       1.1.0.0/22 [90/2713600] via 10.1.14.1, 00:00:27, Virtual-Access2
D       1.1.1.0/24 [90/2713600] via 10.1.14.1, 00:00:27, Virtual-Access2
D       1.1.2.0/24 [90/2713600] via 10.1.14.1, 00:00:27, Virtual-Access2
D       1.1.3.0/24 [90/2713600] via 10.1.14.1, 00:00:27, Virtual-Access2
     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D       10.1.13.3/32 [90/5145600] via 10.1.14.1, 00:00:27, Virtual-Access2
D       10.1.13.0/24 [90/5145600] via 10.1.14.1, 00:00:27, Virtual-Access2
D       10.1.12.0/24 [90/4729856] via 10.1.14.1, 00:00:27, Virtual-Access2
D       10.1.3.0/24 [90/4732416] via 10.1.14.1, 00:00:27, Virtual-Access2
D       10.1.2.0/24 [90/4732416] via 10.1.14.1, 00:00:27, Virtual-Access2
D       10.1.100.0/24 [90/4729856] via 10.1.14.1, 00:00:27, Virtual-Access2
```

```
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.3.0 [90/4857856] via 10.1.14.1, 00:00:29, Virtual-Access2
D       150.1.2.0 [90/4857856] via 10.1.14.1, 00:00:29, Virtual-Access2
```

R4#Show ip eigrp 300 neighbors

```
IP-EIGRP neighbors for process 300
H   Address          Interface     Hold Uptime   SRTT   RTO  Q   Seq
                                   (sec)         (ms)        Cnt Num
0   10.1.14.1        Vi2           12    00:01:16 1311   5000 0   246
```

Task 9

BB2 should be configured such that it advertises EIGRP routes with greater than 110 hops as unreachable.

To see the default setting:

On BB2

BB2#Show ip protocols | Inc EIGRP maximum

```
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
```

To configure the task:

On BB2

```
BB2(config)#router eigrp 200
BB2(config-router)#metric maximum-hops 110
```

The above command will advertise the routes with a hop count higher than 110 as unreachable, the default setting is 100 hops, and it can be increased up to 255.

To verify the configuration:

On BB2

BB2#Show ip protocol | Inc EIGRP maximum

EIGRP maximum hopcount 110
EIGRP maximum metric variance 1

Task 10

The administrative distance of all the routers in AS 200 should be configured as follows:

Internal = 95, External = 138

To see the default setting:

On R2

R2#Show ip protocol | Inc Distance

```
  Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

To change the default values:

On R2

R2(config)#router eigrp 200
R2(config-router)#**distance eigrp 95 138**

On BB2

BB2(config)#router eigrp 200
BB2(config-router)#**distance eigrp 95 138**

The first value after the "distance eigrp" command specifies the administrative distance of Eigrp internal routes and the second value specifies the administrative distance of Eigrp external routes.

To verify the configuration:

On R2

R2#Show ip protocols | Inc Distance

```
  Gateway         Distance      Last Update
Distance: internal 95 external 138
  Gateway         Distance      Last Update
Distance: internal 90 external 170
```

Task 11

BB2 should be configured to use 30 percent of it's F0/0 link for exchanging updates.

On BB2

```
BB2(config)#int f0/0
BB2(config-if)#ip bandwidth-percent eigrp 200 30
```

Task 12

BB1 should be configured to use 15 Mbps of its links bandwidth for exchanging updates. You should NOT use the solution from the previous task to accomplish this task.

On BB1

```
BB1(config)#int f0/0
BB1(config-if)#bandwidth 30000
```

By default Eigrp utilizes 50% of the bandwidth, if you multiply the desired value by two (In this case 30 Mbps) and set the bandwidth of the interface to that number, Eigrp will use half of that number which is the desired value. NOT recommended as the first choice, your first choice should be the solution from the previous task, unless the use of the "IP Bandwidth-percent" command is prohibited.

Task 13

BB1 should be configured to receive routes from R1 and it should not advertise any routes to R1. You should **NOT** use any global configuration command as part of the solution in accomplishing this task.

Note R1 is receiving a single route from BB1:

R1#Sh ip route eigrp 100

```
     150.1.0.0/24 is subnetted, 5 subnets
D       150.1.111.0 [90/25600] via 10.1.111.111, 03:20:34, FastEthernet0/0
```

To configure the task:

On BB1

```
BB1(config)#router eigrp 100
BB1(config-router)#eigrp stub receive-only
```

To verify the configuration:

On R1

R1#Show ip eigrp 100 neighbors

```
IP-EIGRP neighbors for process 100
H   Address         Interface    Hold Uptime   SRTT   RTO  Q   Seq
                                 (sec)         (ms)        Cnt  Num
0   10.1.111.111    Fa0/0        11 00:00:48    1     200   0   13
```

On BB1

BB1#Show ip route eigrp 100

```
     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.1.0 [90/85248] via 10.1.111.1, 00:02:31, FastEthernet0/0
```

Note BB1 only receives routes from R1 and it does NOT advertise any routes to R1, but the neighbor adjacency is maintained.

Task 14

Configure Loopback 1 (151.1.112.112 /24) interface on BB2 and advertise this route in AS 200. This route should appear in the routing table of the routers in this Autonomous System as external.

On BB2

```
BB2(config)#int lo1
BB2(config-if)#ip addr 151.1.112.112 255.255.255.0

BB2(config)#access-list 1 permit 151.1.112.0 0.0.0.255

BB2(config)#Route-map TST permit 10
BB2(config-route-map)#match ip addr 1

BB2(config)#router eigrp 200
BB2(config-router)#redistribute connected route-map TST
```

To verify the configuration:

On R2

R2#Show ip route eigrp 200

```
     151.1.0.0/24 is subnetted, 1 subnets
D       151.1.112.0 [95/130560] via 10.1.112.112, 00:00:57, FastEthernet0/0
     150.1.0.0/24 is subnetted, 4 subnets
D       150.1.112.0 [95/130560] via 10.1.112.112, 00:10:45, FastEthernet0/0
```

Note the reason this network did NOT get injected as an External route is because of the way BB2 is configured, the following reveals Eigrp's configuration of BB2:

BB2#Sh run | S router eigrp

```
router eigrp 200
 redistribute connected route-map TST
 network 0.0.0.0
 metric maximum-hops 110
 metric weights 0 0 0 1 0 0
 distance eigrp 95 138
 no auto-summary
```

The network statement instructs Eigrp to advertise existing and the future configured interfaces in Eigrp AS 200, these routes are internal to Eigrp's AS. Since Internal takes precedence over External routes, the network shows up as an Internal route, to correct this problem, you should reconfigure the network command on BB2 as follows:

```
BB2(config)#router eigrp 200
```

```
BB2(config-router)#NO netw 0.0.0.0
BB2(config-router)#netw 10.1.112.112 0.0.0.0
BB2(config-router)#netw 150.1.112.112 0.0.0.0
```

To verify the configuration:

On R2

```
R2#Sh ip route eigrp 200

     151.1.0.0/24 is subnetted, 1 subnets
D EX    151.1.112.0 [138/130560] via 10.1.112.112, 00:01:16, FastEthernet0/0
     150.1.0.0/24 is subnetted, 4 subnets
D       150.1.112.0 [95/130560] via 10.1.112.112, 00:01:10, FastEthernet0/0
```

Task 15

Configure a static route on BB2 for network 160.1.112.0 /24 using null0 interface as the next hop; this route should be redistributed on BB2. Ensure that existing and future redistributed routes are assigned the following metric:

Bandwidth = 1500
Load = 1
Delay = 20000
Reliability = 255
MTU = 1500

On BB2

```
BB2(config)#ip route 160.1.112.0 255.255.255.0 null0

BB2(config)#router eigrp 200
BB2(config-router)#default-metric 1500 20000 255 1 1500
BB2(config-router)#redistribute static
```

To verify the configuration:

On R2

```
R2#Show ip route eigrp 200
```

```
    160.1.0.0/24 is subnetted, 1 subnets
D EX    160.1.112.0 [138/5122560] via 10.1.112.112, 00:00:59, FastEthernet0/0
    151.1.0.0/24 is subnetted, 1 subnets
D EX    151.1.112.0 [138/130560] via 10.1.112.112, 00:05:31, FastEthernet0/0
    150.1.0.0/24 is subnetted, 4 subnets
D       150.1.112.0 [95/130560] via 10.1.112.112, 00:20:14, FastEthernet0/0
```

Note the default-metric command ONLY affects the static and other redistributed routes but NOT the connected.

Task 16

Configure BB2 such that it ONLY advertises routes that are redistributed and connected networks that are advertised in Eigrp routing protocol. You should NOT use any global configuration command as part of the solution to accomplish this task.

On BB2

```
BB2(config)#Router eigrp 200
BB2(config-router)#eigrp stub connected static
```

To verify the configuration:

On R2

R2#Show ip route eigrp 200

```
        160.1.0.0/24 is subnetted, 1 subnets
D EX    160.1.112.0 [138/5122560] via 10.1.112.112, 00:01:02, FastEthernet0/0
    151.1.0.0/24 is subnetted, 1 subnets
D EX    151.1.112.0 [138/130560] via 10.1.112.112, 00:00:05, FastEthernet0/0
    150.1.0.0/24 is subnetted, 4 subnets
D       150.1.112.0 [95/130560] via 10.1.112.112, 00:01:02, FastEthernet0/0
```

Note the directly connected and ALL redistributed routes are advertised to R2.

Task 17

Configure R5 NOT to log changes in EIGRP neighbor adjacency.

By default Eigrp logs changes in Eigrp neighbor adjacencies. If this is not needed, this feature can be disabled using the following configuration:

On R5

R5(config)#router eigrp 400
R5(config-router)#**NO eigrp log-neighbor-changes**

Task 18

Configure R6 to log neighbor warning messages for the Eigrp 400 and repeat the warning message every 5 minutes.

To enable the logging of "Eigrp neighbor warning messages" you must enter "eigrp log-neighbor-warnings" command under the router eigrp process. This warning message can be repeated based on the number of seconds configured.
By default, neighbor warning messages are logged. If this behavior needs to be changed, then "no eigrp log-neighbor-warning" message must be used.

On R6

R6(config-subif)#router eigrp 400
R6(config-router)#**eigrp log-neighbor-warning 300**

Task 19

Configure R3 to add 50 to the composite metric of all routes received through it's S0/1 interface from router R1.

The following shows the composite metric of all the routes received from R1:

R3#Show ip route eigrp 300

```
     1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       1.1.0.0/22 [90/2297856] via 10.1.100.1, 00:04:28, Serial0/1
D       1.1.1.0/24 [90/2297856] via 10.1.100.1, 00:04:28, Serial0/1
     10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D       10.1.14.4/32 [90/4729856] via 10.1.100.1, 00:04:28, Serial0/1
D       10.1.14.0/24 [90/4729856] via 10.1.100.1, 00:04:28, Serial0/1
D       10.1.12.0/24 [90/2681856] via 10.1.100.1, 00:04:28, Serial0/1
D       10.1.2.0/24 [90/2684416] via 10.1.100.1, 00:04:28, Serial0/1
D       10.1.4.0/24 [90/4732416] via 10.1.100.1, 00:04:28, Serial0/1
     150.1.0.0/24 is subnetted, 3 subnets
D       150.1.4.0 [90/4857856] via 10.1.100.1, 00:04:28, Serial0/1
D       150.1.2.0 [90/2809856] via 10.1.100.1, 00:04:28, Serial0/1
```

To configure Eigrp to add 50 to the existing composite metric:

On R3

Offset-list can be configured to reference an access-list, which references a network/s. If the offset-list references "0" instead of an access-list number, the offset value applies to all the routes received through the specified interface. In this case S0/1.

```
R3(config)#router eigrp 300
R3(config-router)#offset-list 0 in 50 S0/1
```

To verify the configuration:

On R3

R3#Sh ip route eigrp 300

```
     1.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       1.1.0.0/22 [90/2297906] via 10.1.100.1, 00:00:16, Serial0/1
D       1.1.1.0/24 [90/2297906] via 10.1.100.1, 00:00:16, Serial0/1
     10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D       10.1.14.4/32 [90/4729906] via 10.1.100.1, 00:00:16, Serial0/1
D       10.1.14.0/24 [90/4729906] via 10.1.100.1, 00:00:16, Serial0/1
D       10.1.12.0/24 [90/2681906] via 10.1.100.1, 00:00:16, Serial0/1
D       10.1.2.0/24 [90/2684466] via 10.1.100.1, 00:00:16, Serial0/1
D       10.1.4.0/24 [90/4732466] via 10.1.100.1, 00:00:16, Serial0/1
     150.1.0.0/24 is subnetted, 3 subnets
D       150.1.4.0 [90/4857906] via 10.1.100.1, 00:00:16, Serial0/1
D       150.1.2.0 [90/2809906] via 10.1.100.1, 00:00:16, Serial0/1
```

Note a cost of 50 is added to the composite metric of all routes received through S0/1 interface.

Task 20

Configure R4 to perform a mutual redistribution between AS 300 and 400.

On R4

```
R4(config)#router eigrp 400
R4(config-router)#redistribute eigrp 300

R4(config-router)#router eigrp 300
R4(config-router)#redistribute eigrp 400
```

To test the configuration:

On R2

R2#Show ip route eigrp 300

```
     1.0.0.0/22 is subnetted, 1 subnets
D       1.1.0.0 [90/2297856] via 10.1.12.1, 00:36:16, Serial0/0.21
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D       10.1.14.4/32 [90/4729856] via 10.1.12.1, 00:36:16, Serial0/0.21
D       10.1.14.0/24 [90/4729856] via 10.1.12.1, 00:36:16, Serial0/0.21
D       10.1.13.3/32 [90/4729856] via 10.1.12.1, 00:36:16, Serial0/0.21
D       10.1.13.0/24 [90/4729856] via 10.1.12.1, 00:36:16, Serial0/0.21
D       10.1.3.0/24 [90/2684416] via 10.1.12.1, 00:05:09, Serial0/0.21
D       10.1.4.0/24 [90/4732416] via 10.1.12.1, 00:36:16, Serial0/0.21
D EX    10.1.46.0/24 [170/5241856] via 10.1.12.1, 00:00:59, Serial0/0.21
D EX    10.1.45.0/24 [170/8583936] via 10.1.12.1, 00:00:59, Serial0/0.21
D EX    10.1.56.0/24 [170/5244416] via 10.1.12.1, 00:00:59, Serial0/0.21
D       10.1.100.0/24 [90/2681856] via 10.1.12.1, 00:15:34, Serial0/0.21
     150.1.0.0/24 is subnetted, 6 subnets
D EX    150.1.6.0 [170/5369856] via 10.1.12.1, 00:00:59, Serial0/0.21
D EX    150.1.5.0 [170/5372416] via 10.1.12.1, 00:01:00, Serial0/0.21
D       150.1.4.0 [90/4857856] via 10.1.12.1, 00:36:17, Serial0/0.21
D       150.1.3.0 [90/2809856] via 10.1.12.1, 00:05:10, Serial0/0.21
```

Task 21

Configure R2 to inject a default route into AS 200; you should NOT configure any global configuration command as part of the solution to accomplish this task.

On R2

```
R2(config)#int f0/0
R2(config-if)#ip summary-address eigrp 200 0.0.0.0 0.0.0.0
```

To verify the configuration:

On BB2

```
BB2#Show ip route eigrp

D*   0.0.0.0/0 [95/5120] via 10.1.112.2, 00:01:05, FastEthernet0/0
```

Task 22

Configure R1 to perform a mutual redistribution between AS 100 and AS 300, in the future there will be another redistribution point, this router should be configured to prevent feed back routes when the second redistribution point is added.

The routing table of BB1 is checked before the configuration:

```
BB1#Sh ip route eigrp

     150.1.0.0/24 is subnetted, 2 subnets
D       150.1.1.0 [90/85248] via 10.1.111.1, 00:49:49, FastEthernet0/0
```

On R1

```
R1(config)#Route-map 100-300 deny 10
R1(config-route-map)#match tag 300
R1(config)#Route-map 100-300 permit 20
R1(config-route-map)#set tag 100

R1(config)#Route-map 300-100 deny 10
R1(config-route-map)#match tag 100
R1(config)#Route-map 300-100 permit 20
R1(config-route-map)#set tag 300
```

Note tag 100 is set and then denied in the other route-map

The same is performed for tag 300

```
R1(config-route-map)#router eigrp 100
R1(config-router)#redistribute eigrp 300 route-map 300-100

R1(config-router)#Router eigrp 300
R1(config-router)#redistribute eigrp 100 route-map 100-300
```

To verify the configuration:

On BB1

BB1#Show ip route eigrp

```
     1.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D EX    1.1.0.0/24 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    1.1.0.0/22 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    1.1.1.0/24 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    1.1.2.0/24 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    1.1.3.0/24 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D EX    10.1.14.4/32 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.14.0/24 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.13.3/32 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.13.0/24 [170/85248] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.12.0/24 [170/1657856] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.3.0/24 [170/1657856] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.2.0/24 [170/1657856] via 10.1.111.1, 00:00:37, FastEthernet0/0
D EX    10.1.4.0/24 [170/85248] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    10.1.46.0/24 [170/1657856] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    10.1.45.0/24 [170/4999936] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    10.1.56.0/24 [170/1657856] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    10.1.100.0/24 [170/1657856] via 10.1.111.1, 00:00:38, FastEthernet0/0
     150.1.0.0/24 is subnetted, 7 subnets
D EX    150.1.6.0 [170/1657856] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    150.1.5.0 [170/1657856] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    150.1.4.0 [170/85248] via 10.1.111.1, 00:00:38, FastEthernet0/0
D EX    150.1.3.0 [170/1657856] via 10.1.111.1, 00:00:39, FastEthernet0/0
D EX    150.1.2.0 [170/1657856] via 10.1.111.1, 00:00:39, FastEthernet0/0
D       150.1.1.0 [90/85248] via 10.1.111.1, 00:53:22, FastEthernet0/0
```

On R6

R6#Show ip route

```
     1.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
```

```
D EX    1.1.0.0/24 [170/4857856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    1.1.0.0/22 [170/4857856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    1.1.1.0/24 [170/4857856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    1.1.2.0/24 [170/4857856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    1.1.3.0/24 [170/4857856] via 10.1.46.4, 00:16:15, Serial0/0.64
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D EX    10.1.14.1/32 [170/4729856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.14.0/24 [170/4729856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.13.3/32 [170/7289856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.13.0/24 [170/7289856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.12.0/24 [170/5241856] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.3.0/24 [170/5244416] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.2.0/24 [170/5244416] via 10.1.46.4, 00:16:15, Serial0/0.64
D EX    10.1.4.0/24 [170/2172416] via 10.1.46.4, 00:16:16, Serial0/0.64
D       10.1.45.0/24 [90/2172416] via 10.1.56.5, 02:39:08, FastEthernet0/0
D EX    10.1.111.0/24 [170/4732416] via 10.1.46.4, 00:01:52, Serial0/0.64
D EX    10.1.100.0/24 [170/5241856] via 10.1.46.4, 00:16:16, Serial0/0.64
     150.1.0.0/24 is subnetted, 6 subnets
D       150.1.5.0 [90/156160] via 10.1.56.5, 02:39:08, FastEthernet0/0
D EX    150.1.4.0 [170/2297856] via 10.1.46.4, 00:16:16, Serial0/0.64
D EX    150.1.3.0 [170/5369856] via 10.1.46.4, 00:16:16, Serial0/0.64
D EX    150.1.2.0 [170/5369856] via 10.1.46.4, 00:16:16, Serial0/0.64
D EX    150.1.1.0 [170/4857856] via 10.1.46.4, 00:01:54, Serial0/0.64
```

This method is one of the most effective methods used when redistribution between different routing domains occur. In this method, the routes are tagged as they are redistributed and the tags are denied when they are redistributed back.

Task 23

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 2 – Advanced EIGRP Stub Configuration

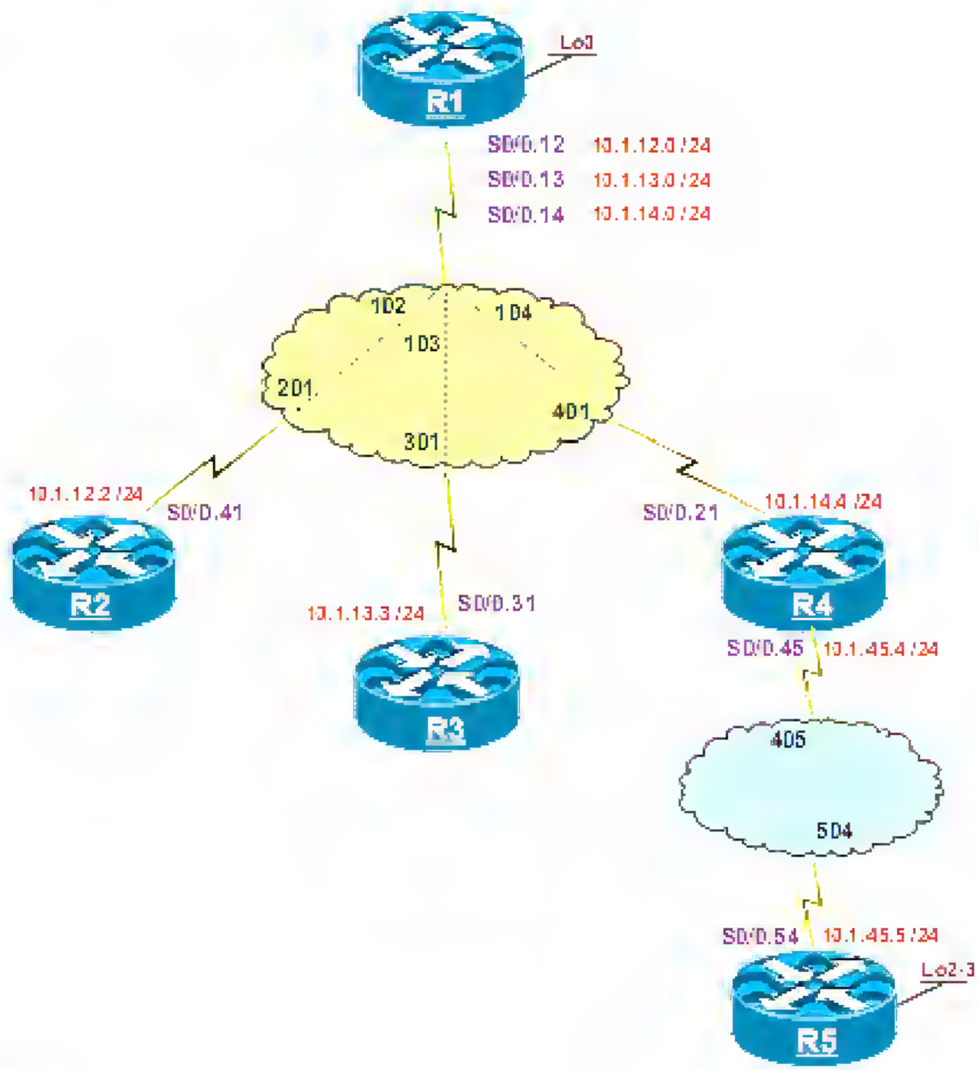

Lab Setup:

- Configure all frame-relay connections in a point-to-point sub-interface manner.
- Use the IP addressing chart below for IP address assignment

IP Addressing:

Router	Interface / IP address	Connecting to router:
R1	S0/0.12 – 10.1.12.1 /24 S0/0.13 – 10.1.13.1 /24 S0/0.14 – 10.1.14.1 /24 Loopback 0 – 1.1.1.1 /24	R2 R3 R4
R2	S0/0.21 – 10.1.12.2 /24	R1
R3	S0/0.31 – 10.1.13.3 /24	R1
R4	S0/0.41 – 10.1.14.4 /24 S0/0.45 – 10.1.45.4 /24	R1 R5
R5	S0/0.54 – 10.1.45.5 /24 Loopback 2 – 2.2.2.2 /24 Loopback 3 – 3.3.3.3 /24	R4

Task 1

Configure OSPF area 0 on the following routers/interfaces; ensure that the loopback interfaces are advertised with their correct mask:

Router	Interface
R4	S0/0.45
R5	S0/0.54 Loopback 2 Loopback 3

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0
```

On R5

```
R5(config)#int lo2
R5(config-if)#ip ospf netw point-to-point

R5(config-if)#int lo3
R5(config-if)#ip ospf netw point-to-point

R5(config)#router ospf 1
```

```
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R5(config-router)#netw 3.3.3.3 0.0.0.0 area 0
```

To verify the configuration:

On R4

```
R4#Show ip route ospf

     2.0.0.0/24 is subnetted, 1 subnets
O       2.2.2.0 [110/65] via 10.1.45.5, 00:01:18, Serial0/0.45
     3.0.0.0/24 is subnetted, 1 subnets
O       3.3.3.0 [110/65] via 10.1.45.5, 00:01:18, Serial0/0.45
```

Task 2

Configure Eigrp 100 on the following routers/interfaces, disable auto summarization:

Router	Interface
R1	S0/0.12 S0/0.13 S0/0.14 Loopback 0
R2	S0/0.21
R3	S0/0.31
R4	S0/0.41

On R1

```
R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#netw 10.1.12.1 0.0.0.0
R1(config-router)#netw 10.1.13.1 0.0.0.0
R1(config-router)#netw 10.1.14.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

On R2

```
R2(config)#router eigrp 100
R2(config-router)#no au
R2(config-router)#netw 10.1.12.2 0.0.0.0
```

On R3

```
R3(config)#router eigrp 100
R3(config-router)#no au
R3(config-router)#netw 10.1.13.3 0.0.0.0
```

On R4

```
R4(config)#router eigrp 100
R4(config-router)#no au
R4(config-router)#netw 10.1.14.4 0.0.0.0
```

To verify the configuration:

On R4

R4#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.14.1, 00:04:29, Serial0/0.41
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.13.0 [90/2681856] via 10.1.14.1, 00:14:24, Serial0/0.41
D       10.1.12.0 [90/2681856] via 10.1.14.1, 00:14:24, Serial0/0.41
```

On R3

R3#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.13.1, 00:03:58, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.13.1, 00:14:48, Serial0/0.31
D       10.1.12.0 [90/2681856] via 10.1.13.1, 00:14:48, Serial0/0.31
```

On R2

R2#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
```

```
D       1.1.1.0 [90/2297856] via 10.1.12.1, 00:03:08, Serial0/0.21
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.12.1, 00:15:16, Serial0/0.21
D       10.1.13.0 [90/2681856] via 10.1.12.1, 00:15:16, Serial0/0.21
```

Task 3

Configure mutual redistribution between OSPF and EIGRP on R4; use a metric of your choice.

On R4

```
R4(config)#router ospf 1
R4(config-router)#redistribute eigrp 100 subnets

R4(config)#router eigrp 100
R4(config-router)#redistribute ospf 1 metric 1 1 1 1 1
```

To verify the configuration:

On R5

R5#Show ip route ospf | Inc O

```
O E2    1.1.1.0 [110/20] via 10.1.45.4, 00:06:00, Serial0/0.54
O E2    10.1.14.0 [110/20] via 10.1.45.4, 00:07:39, Serial0/0.54
O E2    10.1.13.0 [110/20] via 10.1.45.4, 00:07:39, Serial0/0.54
O E2    10.1.12.0 [110/20] via 10.1.45.4, 00:07:39, Serial0/0.54
```

On R1

R1#Show ip route eigrp

```
     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2560512256] via 10.1.14.4, 00:01:00, Serial0/0.14
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0 [170/2560512256] via 10.1.14.4, 00:01:00, Serial0/0.14
     10.0.0.0/24 is subnetted, 4 subnets
D EX    10.1.45.0 [170/2560512256] via 10.1.14.4, 00:01:00, Serial0/0.14
```

On R2

R2#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0    [90/2297856] via 10.1.12.1, 00:07:33, Serial0/0.21
     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2561024256] via 10.1.12.1, 00:01:51, Serial0/0.21
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0    [170/2561024256] via 10.1.12.1, 00:01:51, Serial0/0.21
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.14.0    [90/2681856] via 10.1.12.1, 00:19:41, Serial0/0.21
D       10.1.13.0    [90/2681856] via 10.1.12.1, 00:19:41, Serial0/0.21
D EX    10.1.45.0 [170/2561024256] via 10.1.12.1, 00:01:51, Serial0/0.21
```

On R3

R3#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.13.1, 00:10:45, Serial0/0.31
     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2561024256] via 10.1.13.1, 00:05:02, Serial0/0.31
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0 [170/2561024256] via 10.1.13.1, 00:05:02, Serial0/0.31
     10.0.0.0/24 is subnetted, 4 subnets
D       10.1.14.0 [90/2681856] via 10.1.13.1, 00:21:34, Serial0/0.31
D       10.1.12.0 [90/2681856] via 10.1.13.1, 00:21:34, Serial0/0.31
D EX    10.1.45.0 [170/2561024256] via 10.1.13.1, 00:05:02, Serial0/0.31
```

Task 4

Configure "Eigrp stub" on R1 such that it ONLY advertises it's directly connected interfaces that are advertised with a "network" command to its Eigrp neighbors.

On R1

R1(config)#router eigrp 100
R1(config-router)#**eigrp stub connected**

To verify the configuration:

On R2

R2#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.12.1, 00:02:01, Serial0/0.21
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.12.1, 00:02:01, Serial0/0.21
D       10.1.13.0 [90/2681856] via 10.1.12.1, 00:02:01, Serial0/0.21
```

On R3

R3#Show ip route eigrp

```
     1.0.0.0/24 is subnetted, 1 subnets
D       1.1.1.0 [90/2297856] via 10.1.13.1, 00:02:41, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.14.0 [90/2681856] via 10.1.13.1, 00:02:41, Serial0/0.31
D       10.1.12.0 [90/2681856] via 10.1.13.1, 00:02:41, Serial0/0.31
```

On R1

R1#Show ip route eigrp

```
     2.0.0.0/24 is subnetted, 1 subnets
D EX    2.2.2.0 [170/2560512256] via 10.1.14.4, 00:03:25, Serial0/0.14
     3.0.0.0/24 is subnetted, 1 subnets
D EX    3.3.3.0 [170/2560512256] via 10.1.14.4, 00:03:25, Serial0/0.14
     10.0.0.0/24 is subnetted, 4 subnets
D EX    10.1.45.0 [170/2560512256] via 10.1.14.4, 00:03:25, Serial0/0.14
```

Note R2 and R3 do NOT get the redistributed routes, because of R1's "stub connected" configuration.

Task 5

Configure R1 such that routers R2 and R3 have networks 2.2.2.0 /24 and 3.3.3.0 /24 in their routing table. DO NOT remove the "Eigrp stub connected" configuration from R1 to accomplish this task.

To accomplish this task a "leak-map" is referenced in the "Eigrp stub connected"

command. The leak-map references a route-map called "TST", the route-map references an access-list. Any IP address/es that are permitted in the access-list is leaked along the connected networks.

On R1

R1(config)#access-list 1 permit 2.2.2.0 0.0.0.255
R1(config)#access-list 1 permit 3.3.3.0 0.0.0.255

R1(config)#route-map TST permit 10
R1(config-route-map)#**match ip addr** 1

R1(config)#router eigrp 100
R1(config-router)#**eigrp stub connected leak-map TST**

To verify the configuration:

On R2

R2#Show ip route eigrp | Inc EX

D EX 2.2.2.0 [170/2561024256] via 10.1.12.1, 00:04:29, Serial0/0.21
D EX 3.3.3.0 [170/2561024256] via 10.1.12.1, 00:04:29, Serial0/0.21

On R3

R3#Show ip route eigrp | Inc EX

D EX 2.2.2.0 [170/2561024256] via 10.1.13.1, 00:00:13, Serial0/0.31
D EX 3.3.3.0 [170/2561024256] via 10.1.13.1, 00:00:13, Serial0/0.31

Note both R2 and R3 have both networks 2.2.2.0 /24 and 3.3.3.0 /24 in their routing table.

Task 6

Re-configure R1 such that R2 gets network 2.2.2.0 /24 and R3 gets network 3.3.3.0 /24 ONLY. DO NOT remove the "Eigrp stub connected" configuration from R1 to accomplish this task.

The access-list and the route-map TST should be removed before proceeding further:

On R1

R1(config)#**NO** access-list 1
R1(config)#**NO** route-map TST

We should identify the two networks using two access-lists, in this case access-list 2 permits network 2.2.2.0 /24 and access-list 3 permits network 3.3.3.0 /24:

R1(config)#access-list **2** permit 2.2.2.0 0.0.0.255
R1(config)#access-list **3** permit 3.3.3.0 0.0.0.255

The next step is to configure a new route-map as follows:

R1(config)#route-map **TST permit 10**
R1(config-route-map)#match ip addr **2**
R1(config-route-map)#match inter S0/0.12

R1(config-route-map)#route-map **TST permit 20**
R1(config-route-map)#match ip addr **3**
R1(config-route-map)#match inter S0/0.13

Note the route-map is already referenced by the leak-map.

To verify the configuration:

On R2

R2#Show ip route eigrp | Inc EX

D EX **2.2.2.0** [170/2561024256] via 10.1.12.1, 00:03:16, Serial0/0.21

On R3

R3#Show ip route eigrp | Inc EX

D EX **3.3.3.0** [170/2561024256] via 10.1.13.1, 00:03:43, Serial0/0.31

Task 7

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 3 – Eigrp & Default-information

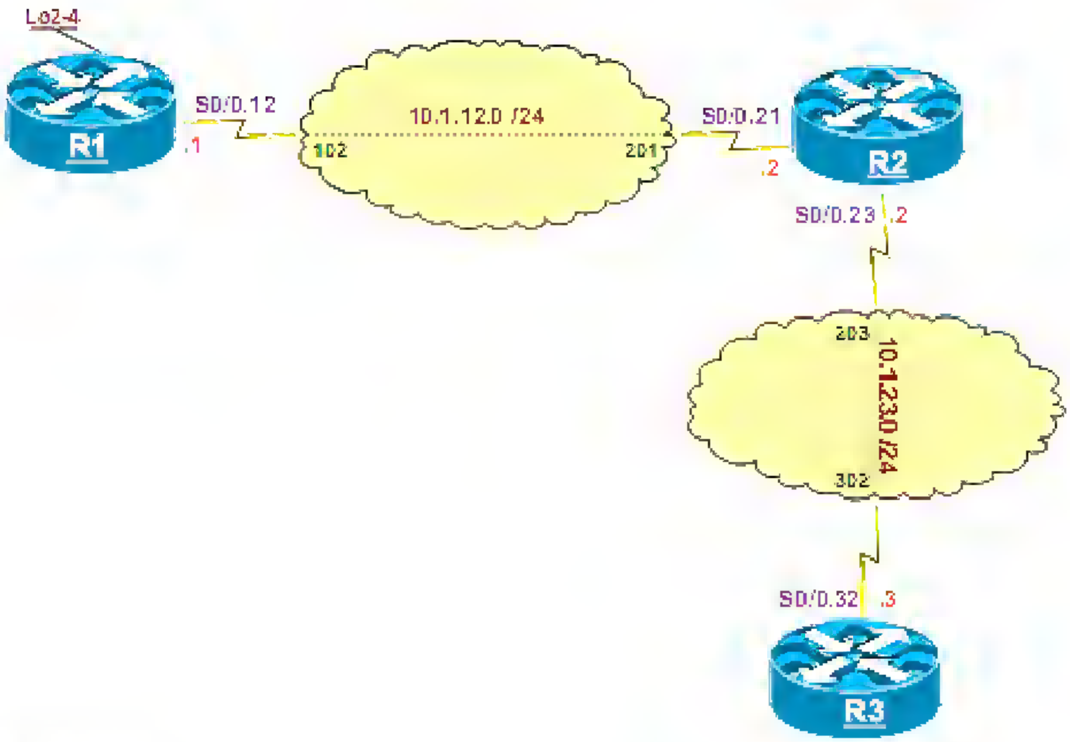

Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Use the IP addressing and DLCI chart below.

IP addressing:

Router	Interface / IP address	DLCI assignment
R1	S0/0.12 = 10.1.12.1 /24 Loopback2 = 2.2.2.2 /8 Loopback3 = 3.3.3.3 /8 Loopback4 = 4.4.4.4 /8	102
R2	S0/0.21 = 10.1.12.2 /24 S0/0.23 = 10.1.23.2 /24	201 203
R3	S0/0.32 = 10.1.23.3 /24	302

Task 1

Configure Eigrp on all routers and advertise their directly connected interfaces in AS 100. R1 should NOT advertise network 4.0.0.0 /8 in this routing protocol.

On R1

R1(config)#Router eigrp 100
R1(config-router)#no au
R1(config-router)#netw 2.2.2.2 0.0.0.0
R1(config-router)#netw 3.3.3.3 0.0.0.0
R1(config-router)#netw 10.1.12.1 0.0.0.0

On R2

R2(config)#Router eigrp 100
R2(config-router)#no au
R2(config-router)#network 10.1.12.2 0.0.0.0
R2(config-router)#network 10.1.23.2 0.0.0.0

On R3

R3(config)#Router eigrp 100
R3(config-router)#no au
R3(config-router)#network 10.1.23.3 0.0.0.0

To verify the configuration:

On R1

R1#Show ip route eigrp

```
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.23.0 [90/2681856] via 10.1.12.2, 00:01:55, Serial0/0.12
```

On R2

R2#Show ip route eigrp

```
D    2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:02:21, Serial0/0.21
D    3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:02:21, Serial0/0.21
```

On R3

R3#Show ip route eigrp

```
D    2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:57, Serial0/0.32
D    3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:57, Serial0/0.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:57, Serial0/0.32
```

Task 2

Configure R1 such that R2 and R3 use network 2.0.0.0 /8 as candidate default.

On R1

R1(config)#**ip default-network 2.0.0.0**

To verify the configuration:

On R1

R1#Sh ip route 2.0.0.0

```
Routing entry for 2.0.0.0/8
  Known via "connected", distance 0, metric 0 (connected, via interface), candidate default path
  Redistributing via eigrp 100
  Routing Descriptor Blocks:
  * directly connected, via Loopback2
      Route metric is 0, traffic share count is 1
```

R1#Show ip route | B Gate

```
Gateway of last resort is not set

C*   2.0.0.0/8 is directly connected, Loopback2
C    3.0.0.0/8 is directly connected, Loopback3
C    4.0.0.0/8 is directly connected, Loopback4
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
D       10.1.23.0 [90/2681856] via 10.1.12.2, 00:05:13, Serial0/0.12
```

On R2

R2#Show ip route eigrp

D* 2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:00:41, Serial0/1.21
D 3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:07:55, Serial0/1.21

Note the asterisk reveals that R2 is using that network as the candidate default.

On R3

R3#Sh ip route eigrp

D* 2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:00:23, Serial0/0.32
D 3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:39, Serial0/0.32
10.0.0.0/24 is subnetted, 2 subnets
D 10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:39, Serial0/0.32

To test the configuration:

On R2

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/26/60 ms

To verify the configuration on R3:

R3#Show ip route eigrp

D* 2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:04:24, Serial0/1.32
D 3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:10:39, Serial0/1.32
10.0.0.0/24 is subnetted, 2 subnets
D 10.1.12.0 [90/2681856] via 10.1.23.2, 00:10:39, Serial0/1.32

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/32/92 ms

Task 3

Configure R2 such that R3 does NOT use network 2.0.0.0 /8 as the candidate default, R3 should still have network 2.0.0.0 /8 in it's routing table.

On R2

R2(config)#router eigrp 100
R2(config-router)#NO **default-information allowed** out

R2#Clear ip eigrp neighbor

Note the "NO default-information allowed out" disables the redistribution of default route, meaning that R3 will no longer use network 2.0.0.0 /8 as its candidate default, but it will still have that network in its routing table.

To verify the configuration:

On R2

R2#Show ip route eigrp

```
D*   2.0.0.0/8 [90/2297856] via 10.1.12.1, 00:00:25, Serial0/1.21
D    3.0.0.0/8 [90/2297856] via 10.1.12.1, 00:00:25, Serial0/1.21
```

On R3

R3#Show ip route eigrp

```
D    2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:04, Serial0/1.32
D    3.0.0.0/8 [90/2809856] via 10.1.23.2, 00:01:04, Serial0/1.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:06, Serial0/1.32
```

To test the configuration:

On R3

R3#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
.....
```

Success rate is 0 percent (0/5)

Task 4

Remove the command that was entered in the previous task on R2, and configure R3 to accomplish the same task, if this configuration is performed correctly, R3 should NOT use network 2.0.0.0 /8 as the candidate default.

To remove the command from R2:

On R2

```
R2(config)#router eigrp 100
R2(config-router)#default-information allowed out

R2#Clear ip eigrp neighbor
```

To verify R2's configuration:

On R2

```
R2#Sh run | S router eigrp

router eigrp 100
 network 10.1.12.2 0.0.0.0
 network 10.1.23.2 0.0.0.0
 no auto-summary
```

To verify the configuration on R3

On R3

```
R3#Show ip route eigrp

D*   2.0.0.0/8 [90/2809856] via 10.1.23.2, 00:00:58, Serial0/1.32
D    3.0.0.0/8  [90/2809856] via 10.1.23.2, 00:00:58, Serial0/1.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:01:00, Serial0/1.32
```

Note R3 uses network 2.0.0.0 /8 as candidate default. To configure R3 to

accomplish the same task:

On R3

R3(config)#router eigrp 100
R3(config-router)#NO default-information allowed in

R3#Clear ip eigrp neighbor

Note from R3's perspective it should disable the redistribution of the default route inbound, therefore, the direction of the command is configured inbound.

Note if R2 enters the "no default-information allowed in", then R3 will not receive is either.

To test and verify the configuration:

On R3

R3#Show ip route eigrp

```
D    2.0.0.0/8     [90/2809856] via 10.1.23.2, 00:02:38, Serial0/1.32
D    3.0.0.0/8     [90/2809856] via 10.1.23.2, 00:02:38, Serial0/1.32
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/2681856] via 10.1.23.2, 00:02:38, Serial0/1.32
```

Task 5

Reconfigure the routers based on the following topology and IP addressing:

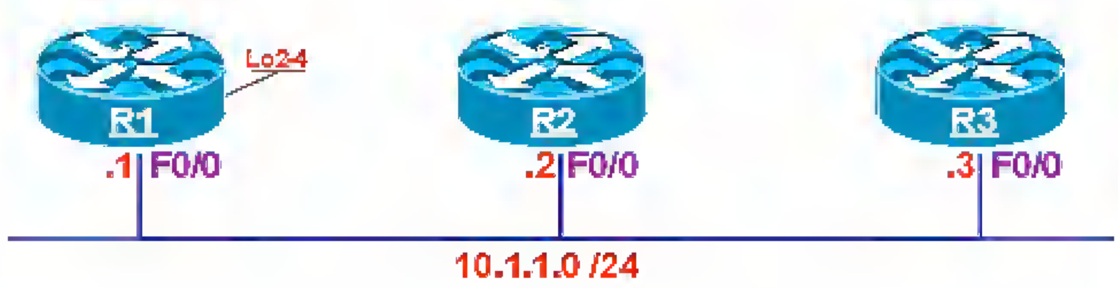

Lab Setup:

- Configure F0/0 interface of the routers in VLAN 100.
- Use the IP addressing chart below.

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 Loopback2 = 2.2.2.2 /8 Loopback3 = 3.3.3.3 /8 Loopback4 = 4.4.4.4 /8
R2	**F0/0 = 10.1.1.2 /24**
R3	F0/0 = 10.1.1.3 /24

Task 6

Configure Eigrp 100 on the routers and advertise their directly connected interfaces in AS 100. R1 should NOT advertise network 4.0.0.0 /8 in this routing protocol.

On R1

```
R1(config-if)#router eigrp 100
R1(config-router)#no au
R1(config-router)#network 2.0.0.0
R1(config-router)#Network 3.0.0.0
R1(config-router)#Network 10.1.1.1 0.0.0.0
```

On R2

```
R2(config)#Router eigrp 100
R2(config-router)#no au
R2(config-router)#network 10.1.1.2 0.0.0.0
```

On R3

```
R3(config)#Router eigrp 100
```

R3(config-router)#no au
R3(config-router)#network 10.1.1.3 0.0.0.0

To verify the configuration

On R2

R2#Show ip route eigrp

D 2.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:21, Ethernet0/0
D 3.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:21, Ethernet0/0

On R3

R3#Show ip route eigrp

D 2.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:06, Ethernet0/0
D 3.0.0.0/8 [90/409600] via 10.1.1.1, 00:15:06, Ethernet0/0

Note R1 should not have any Eigrp routes in its routing table.

Task 7

Configure R1 to advertise Network 2.0.0.0 /8 and Network 3.0.0.0 /8 as candidate default in this routing domain.

On R1

R1(config)#**ip default-network 2.0.0.0**
R1(config)#**ip default-network 3.0.0.0**

To verify the configuration:

On R2

R2#Show ip route eigrp

D* 2.0.0.0/8 [90/409600] via 10.1.1.1, 00:02:55, Ethernet0/0
D* 3.0.0.0/8 [90/409600] via 10.1.1.1, 00:02:52, Ethernet0/0

On R3

R3#Show ip route eigrp

D* 2.0.0.0/8 [90/409600] via 10.1.1.1, 00:03:56, Ethernet0/0
D* 3.0.0.0/8 [90/409600] via 10.1.1.1, 00:03:52, Ethernet0/0

Note both R2 and R3 use networks 2.0.0.0 /8 and 3.0.0.0 /8 as their candidate default.

Task 8

Configure R2 and R3 such that R2 uses network 2.0.0.0 /8 and R3 uses network 3.0.0.0 /8 as their candidate default.

On R2

To configure this task, an access-list is written to identify the network (Network 2.0.0.0 /8 in this case), then, the access-list is referenced in the "Default-information allowed in" command, which tells the router that ONLY the network that is permitted in the access-list should be used as candidate default.

R2(config)#access-list 2 permit 2.0.0.0

R2(config)#router eigrp 100
R2(config-router)#default-information allowed in 2

R2#clc ip eigrp neigh

R2#Show ip route eigrp

D* 2.0.0.0/8 [90/409600] via 10.1.1.1, 00:01:41, FastEthernet0/0
D 3.0.0.0/8 [90/409600] via 10.1.1.1, 00:01:41, FastEthernet0/0

The following shows the configuration of R3:

On R3

R3(config)#access-list 3 permit 3.0.0.0

R3(config)#Router eigrp 100
R3(config-router)#default-information allowed in 3

R3#Cle ip eigrp neighbor

To verify the configuration:

On R3

R3#Show ip route eigrp

D 2.0.0.0/8 [90/409600] via 10.1.1.1, 00:01:01, Ethernet0/0
D* 3.0.0.0/8 [90/409600] via 10.1.1.1, 00:01:01, Ethernet0/0

To test the configuration:

On R2

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/11/24 ms

On R3

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/13/24 ms

Task 9

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 4 – Eigrp Filtering

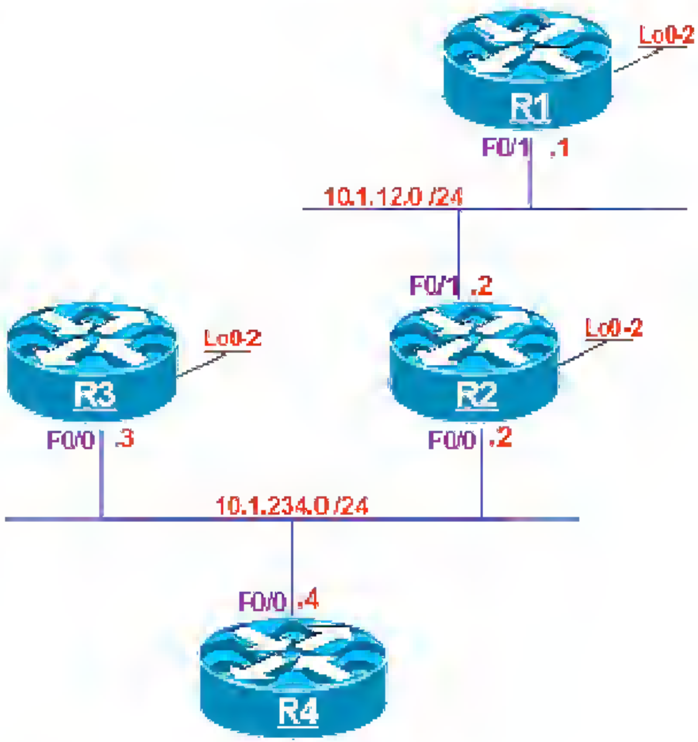

Lab Setup:

- Configure the F0/1 interface of R1 and R2 in VLAN 12
- Configure the F0/0 interface of R2, R3 and R4 in VLAN 234
- Use the following IP addressing chart for IP assignment:

IP Addressing chart:

Router	Interface / IP address
R1	F0/1 = 10.1.12.1 /24 Lo0 = 1.1.1.1 /8 Lo1 = 11.1.1.1 /8 Lo2 = 111.1.1.1 /8
R2	F0/1 = 10.1.12.2 /24 F0/0 = 10.1.234.2 /24

	Lo0 = 2.2.2.2 /8 Lo1 = 200.1.1.1 /24 Lo2 = 200.2.2.2 /24
R3	F0/0 = 10.1.234.3 /24 Lo0 = 3.3.3.3 /8 Lo1 = 200.1.1.1 /24 Lo2 = 200.2.2.2 /24
R4	F0/0 = 10.1.234.4 /24

Task 1

Configure EIGRP 100 on all routers and advertise their directly connected links. You should disable auto summarization on these routers

On R1

```
R1(config)#router eigrp 100
R1(config-router)#network 1.1.1.1 0.0.0.0
R1(config-router)#network 10.1.12.1 0.0.0.0
R1(config-router)#network 11.1.1.1 0.0.0.0
R1(config-router)#network 111.1.1.1 0.0.0.0
R1(config-router)#no auto-summary
```

On R2

```
R2(config)#router eigrp 100
R2(config-router)#network 10.1.12.2 0.0.0.0
R2(config-router)#network 10.1.234.2 0.0.0.0
R2(config-router)#network 200.1.1.1 0.0.0.0
R2(config-router)#network 200.2.2.2 0.0.0.0
R2(config-router)#network 2.2.2.2 0.0.0.0
R2(config-router)#no auto-summary
```

On R3

```
R3(config)#router eigrp 100
R3(config-router)#network 10.1.234.3 0.0.0.0
R3(config-router)#network 200.1.1.1 0.0.0.0
R3(config-router)#network 200.2.2.2 0.0.0.0
R3(config-router)#network 3.3.3.3 0.0.0.0
R3(config-router)#no auto-summary
```

On R4

```
R4(config)#router eigrp 100
R4(config-router)#network 10.1.234.4 0.0.0.0
R4(config-router)#no auto-summary
```

To verify the configuration:

On R4

R4#Show ip route Eigrp

```
D    1.0.0.0/8 [90/158720] via 10.1.234.2, 00:00:17, FastEthernet0/0
D    2.0.0.0/8 [90/156160] via 10.1.234.2, 00:00:17, FastEthernet0/0
D    3.0.0.0/8 [90/156160] via 10.1.234.3, 00:00:17, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.3, 00:00:17, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:17, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:00:17, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:17, FastEthernet0/0
D    111.0.0.0/8 [90/158720] via 10.1.234.2, 00:00:17, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0 [90/30720] via 10.1.234.2, 00:00:17, FastEthernet0/0
D    11.0.0.0/8   [90/158720] via 10.1.234.2, 00:00:17, FastEthernet0/0
```

Task 2

Configure R4 such that it filters existing (1.0.0.0/8, 11.0.0.0/8 and 111.0.0.0/8) and future network behind R1. DO NOT use "distribute-list" or "route-map" to accomplish this task

By default, Eigrp will discard routes that have a hop count of 101 or more, this behavior can be utilized to accomplish this task, as follows:

On R4

R4#Sh ip eigrp topology 11.0.0.0 255.0.0.0 | Inc Hop

```
        Hop count is 2
```

R4#Sh ip eigrp topology 1.0.0.0 255.0.0.0 | Inc Hop

```
        Hop count is 2
```

R4#Sh ip eigrp topology 111.0.0.0 255.0.0.0 | Inc Hop

Hop count is **2**

Note the routes behind R1 have a hop count of 2, whereas, the other routes advertised in this topology have a hop count of 1:

R4#Sh ip eigrp topology 2.0.0.0 255.0.0.0 | Inc Hop

Hop count is **1**

R4#Sh ip eigrp topology 3.0.0.0 255.0.0.0 | Inc Hop

Hop count is **1**

Therefore, we should reject routes that have a hop count greater than 1:

On R4

R4(config)#router eigrp 100
R4(config-router)#**metric maximum-hops** 1

Note when the above command is entered, the following message should be received, this is because the policy for Eigrp is changed from 100 (Default hop count) to ONLY 1:

*%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.2 (FastEthernet0/0) is down: **Max hopcount changed***
*%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.3 (FastEthernet0/0) is down: **Max hopcount changed***
*%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.2 (FastEthernet0/0) is up: **new adjacency***
*%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.234.3 (FastEthernet0/0) is up: **new adjacency***

To verify the configuration:

On R4

R4#Show ip route eigrp

```
D    2.0.0.0/8      [90/156160] via 10.1.234.2, 00:01:48, FastEthernet0/0
D    3.0.0.0/8      [90/156160] via 10.1.234.3, 00:01:48, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.3, 00:01:48, FastEthernet0/0
```

```
                   [90/156160] via 10.1.234.2, 00:01:48, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:01:48, FastEthernet0/0
                   [90/156160] via 10.1.234.2, 00:01:48, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0  [90/30720] via 10.1.234.2, 00:01:48, FastEthernet0/0
```

Task 3

Configure R4 such that it ONLY takes R2 to reach Network 200.1.1.0 /24.

On R4

Note in this case an extended access-list can be used to filter a prefix from a given route-source; in the following extended access-list, the source address in the ACL references the advertising neighbor, whereas, the destination address in the ACL references the actual Network:

```
R4(config)#access-list 100 deny ip host 10.1.234.3 host 200.1.1.0
R4(config)#access-list 100 permit ip any any

R4(config)#router eigrp 100
R4(config-router)#distribute-list 100 in F0/0
```

To verify the configuration:

On R4

R4#Sh ip route eigrp

```
D    2.0.0.0/8     [90/156160] via 10.1.234.2, 00:01:20, FastEthernet0/0
D    3.0.0.0/8     [90/156160] via 10.1.234.3, 00:01:20, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.2, 00:01:20, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:01:20, FastEthernet0/0
                   [90/156160] via 10.1.234.2, 00:01:20, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0  [90/30720] via 10.1.234.2, 00:01:20, FastEthernet0/0
```

Task 4

Re-configure the solution in task 3 such that if R2 is down R4 can reach network 200.1.1.0/24 through R3.

In this case the distance can be manipulated to accomplish this task, as follows:

To remove the commands from the previous step:

R4(config)#router eigrp 100
R4(config-router)#No **distribute-list 100 in F0/0**

R4(config)#No access-list 100

The next step is to configure the new policy:

Step 1:
Configure an access-list to identify the network:

R4(config)#**access-list** 1 **permit 200.1.1.0 0.0.0.255**

Step 2:
Utilizing the distance command, the AD for network 200.1.1.0/24 ONLY through R3 is set higher than the default AD of 90:

R4(config-router)#**distance** 91 10.1.234.3 0.0.0.0 1

The above command sets the AD to 91 through R3 for networks identified in access-list 1. The following command resets the neighbors (This is done to speed up the process, on some IOS versions, it is done automatically:

R4#Cle ip eigrp neighbor

To verify the configuration:

On R4

R4#Sh ip route eigrp

```
D    2.0.0.0/8       [90/156160] via 10.1.234.2, 00:00:12, FastEthernet0/0
D    3.0.0.0/8       [90/156160] via 10.1.234.3, 00:00:12, FastEthernet0/0
D    200.1.1.0/24 [90/156160] via 10.1.234.2, 00:00:12, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:00:12, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:00:12, FastEthernet0/0
```

```
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0  [90/30720] via 10.1.234.2, 00:00:12, FastEthernet0/0
```

To test the configuration:

On R2

```
R2(config)#int lo1
R2(config-if)#Shut
```

To verify the test:

On R4

```
R4#Sh ip route eigrp

D    2.0.0.0/8      [90/156160] via 10.1.234.2, 00:01:05, FastEthernet0/0
D    3.0.0.0/8      [90/156160] via 10.1.234.3, 00:01:05, FastEthernet0/0
D    200.1.1.0/24 [91/156160] via 10.1.234.3, 00:00:08, FastEthernet0/0
D    200.2.2.0/24 [90/156160] via 10.1.234.3, 00:01:05, FastEthernet0/0
                  [90/156160] via 10.1.234.2, 00:01:05, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0  [90/30720] via 10.1.234.2, 00:01:05, FastEthernet0/0
```

Task 5

Filter network 2.0.0.0/8 on R4; DO NOT use the "distribute-list" command to accomplish this task.

Once again the distance command can be used to accomplish this task, the difference between the solution used in this task and the solution used in the previous task is that the AD is set to a value that is unreachable (255).

On R4

```
R4(config)#access-list 2 permit 2.0.0.0

R4(config)#router eigrp 100
R4(config-router)#distance 255 10.1.234.2 0.0.0.0 2
```

To verify the configuration:

On R4

R4#Show ip route eigrp

```
D    3.0.0.0/8      [90/156160] via 10.1.234.3, 00:02:16, FastEthernet0/0
D    200.1.1.0/24   [90/156160] via 10.1.234.2, 00:02:16, FastEthernet0/0
D    200.2.2.0/24   [90/156160] via 10.1.234.3, 00:02:16, FastEthernet0/0
                    [90/156160] via 10.1.234.2, 00:02:16, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.12.0   [90/30720] via 10.1.234.2, 00:02:16, FastEthernet0/0
```

Task 6

Erase the startup config and reload the routers before proceeding to the next task.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

OSPF

Lab 1 – Advertising Networks

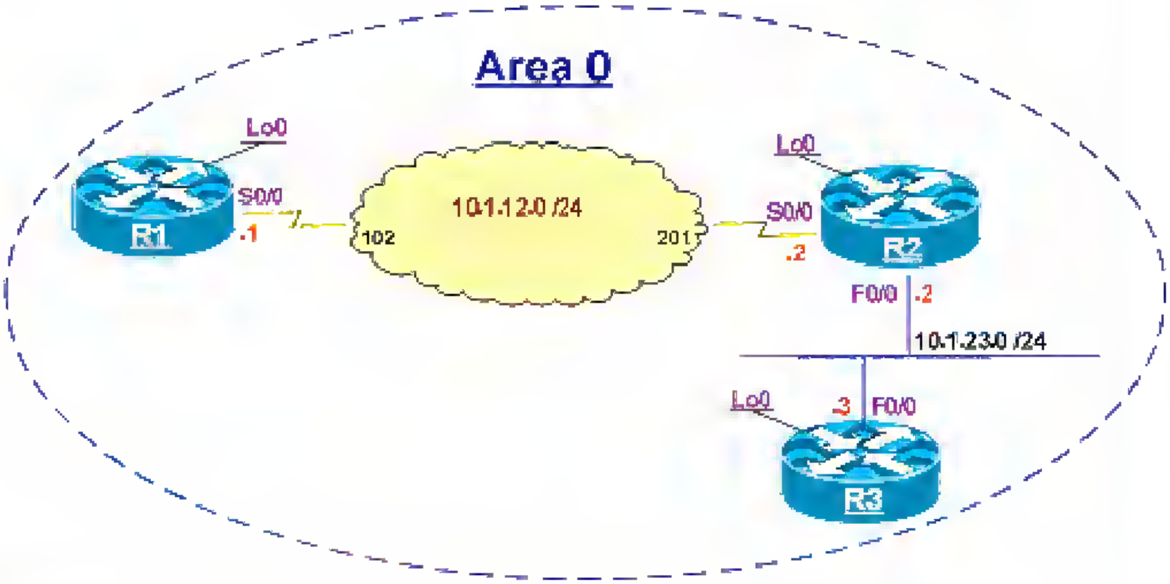

Lab Setup:

- Configure the frame-relay connection using the S0/0 interface
- Configure the F0/0 interface of R2 and R3 in VLAN 23
- Use the following IP addressing chart for IP address assignment

IP addressing:

Router	Interface / IP addressing
R1	S0/0 = 10.1.12.1 /24 Lo0 = 1.1.1.1 /8
R2	S0/0 = 10.1.12.2 /24 F0/0 = 10.1.23.2 /24 10.2.2.2 /24 Secondary Lo0 = 2.2.2.2 /8
R3	F0/0 = 10.1.23.3 /24 10.3.3.3 /24 Secondary Lo0 = 3.3.3.3 /8

Task 1

Configure OSPF on these routers and run every interface (This includes the secondary interfaces) of these routers in Area 0; do not use "Network" command to accomplish this task. The loopback interface/s should be advertised with their correct mask. There should NOT be a DR election on the Frame-relay network, do NOT use Point-to-Multipoint network type.

On R1

```
R1(config)#int S0/0
R1(config-if)#ip ospf network point-to-point
R1(config-if)#ip ospf 1 area 0

R1(config)#int lo0
R1(config-if)#ip ospf network point-to-point
R1(config-if)#ip ospf 1 area 0
```

On R2

```
R2(config)#int Lo0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#ip ospf 1 area 0

R2(config)#int F0/0
R2(config-if)#ip ospf 1 area 0

R2(config)#int S0/0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#ip ospf 1 area 0
```

On R3

```
R3(config)#int Lo0
R3(config-if)#ip ospf 1 area 0
R3(config-if)#ip ospf network point-to-point

R3(config)#int F0/0
R3(config-if)#ip ospf 1 area 0
```

To verify the configuration:

On R1

R1#Show ip route ospf 1 | inc O

```
O       2.0.0.0/8  [110/65] via 10.1.12.2, 00:05:58, Serial0/0
O       3.0.0.0/8  [110/75] via 10.1.12.2, 00:05:58, Serial0/0
O       10.3.3.0   [110/84] via 10.1.12.2, 00:05:58, Serial0/0
O       10.2.2.0   [110/74] via 10.1.12.2, 00:05:58, Serial0/0
O       10.1.23.0 [110/74] via 10.1.12.2, 00:05:58, Serial0/0
```

Note the secondary IP addresses are also advertised.

Task 2

Configure R2 and R3 such that the secondary IP addresses are NOT advertised; do NOT use Access-list, Prefix-lists or filtering of any type and minimum number of commands should be used to accomplish this task.

On R2 and R3

R2(config)#int F0/0
R2(config-if)#**ip ospf 1 area 0** secondaries none

To verify the configuration:

On R1

R1#Show ip route ospf 1 | inc O

```
O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:01:00, Serial0/0
O    3.0.0.0/8 [110/75] via 10.1.12.2, 00:01:00, Serial0/0
O       10.1.23.0 [110/74] via 10.1.12.2, 00:01:00, Serial0/0
```

Note the secondary Prefixes are no longer advertised.

Task 3

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 2
Optimization of OSPF and Adjusting Timers

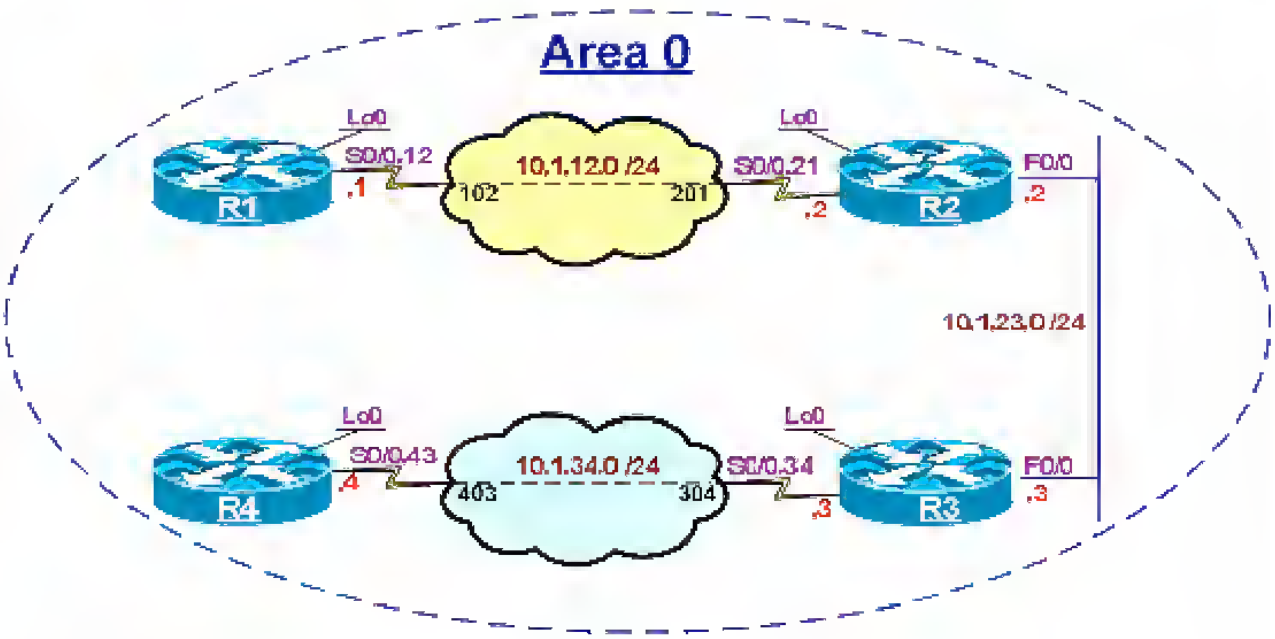

Lab Setup:

- R2 and R3's F0/0 interface should be configured in VLAN 23.
- R1 and R2, R3 and R4 should be configured in a frame-relay point-to-point manner.

IP Addressing

Router	Interface	IP address
R1	Lo0 F/R interface	1.1.1.1 /8 10.1.12.1 /24
R2	Lo0 F/R interface F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24
R3	Lo0 F/R interface F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24
R4	Lo0 F/R interface	4.4.4.4 /8 10.1.34.4 /24

Task 1

Configure OSPF on all routers and advertise their directly connected networks in area 0.

On R4

```
R4(config)#Router ospf 1
R4(config-router)#net 0.0.0.0 0.0.0.0 are 0
```

On R3

```
R3(config-if)#Router ospf 1
R3(config-router)#network 0.0.0.0 0.0.0.0 are 0
```

On R2

```
R2(config-if)#Router ospf 1
R2(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

On R1

```
R1(config-if)#Router ospf 1
R1(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

To verify the configuration:

On R1

R1#Show ip route ospf

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:04, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/66] via 10.1.12.2, 00:00:04, Serial0/0.12
     4.0.0.0/32 is subnetted, 1 subnets
O       4.4.4.4 [110/130] via 10.1.12.2, 00:00:04, Serial0/0.12
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.23.0 [110/65] via 10.1.12.2, 00:00:04, Serial0/0.12
O       10.1.34.0 [110/129] via 10.1.12.2, 00:00:04, Serial0/0.12
```

Task 2

R4 is getting flooded with LSA type 6 packets; ensure that R4 does not generate a syslog message for these packets.

Cisco routers do NOT support LSA type 6 packets and each time an OSPF router receives an MOSPF (LSA type 6) packet it sends a syslog message. If the routers receive many LSA type 6 packets they will generate a large number of syslog messages. This feature should be disabled to prevent this from occurring.

On R4

```
R4(config)#router ospf 1
R4(config-router)#ignore lsa mospf
```

Task 3

To ensure fast detection of a neighbor being down, configure R2 and R3 to send their hellos four times a second with a hold time of one second.

On R2

```
R2(config)#int f0/0
R2(config-if)#ip ospf dead-interval minimal hello-multiplier 4
```

On R3

```
R3(config-if)#int f0/0
R3(config-if)#ip ospf dead-interval minimal hello-multiplier 4
```

The dead interval is advertised in OSPF hello packets. The values of this parameter must be the same for two routers in order for them to form a neighbor adjacency. By specifying the "minimal" and "hello-multiplier" keywords with a multiplier value, you are enabling OSPF fast hello packets. The "minimal" keyword sets the dead interval to 1 second and the "hello-multiplier" value sets the number of hello packets sent during that 1 second.

Task 4

Ensure that these routers lookup DNS names for use in all OSPF show commands, test this task to ensure proper operation. Since there are no DNS servers in this lab you should use the routers for that purpose.

To test the OSPF "Show" commands before implementing this feature, enter the following:

Show ip ospf database router

```
        OSPF Router with ID (1.1.1.1) (Process ID 1)

            Router Link States (Area 0)

  LS age: 1575
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000002
  Checksum: 0x1D3F
  Length: 60
  Number of Links: 3
```

Note the router-id is displayed

On All Routers

```
(config)#ip ospf name-lookup
(config)#ip host R1 1.1.1.1
(config)#ip host R2 2.2.2.2
(config)#ip host R3 3.3.3.3
(config)#ip host R4 4.4.4.4
```

Show ip ospf database router

```
        OSPF Router with ID (1.1.1.1) (Process ID 1)

            Router Link States (Area 0)

  LS age: 1651
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: R1
```

Note the change, its replaced by the name configured in the "IP host" command.

```
  LS Seq Number: 80000002
Checksum: 0x1D3F
  Length: 60
  Number of Links: 3
```

Task 5

Configure R2 such that if it does not receive an acknowledgment from R3 for a given LSA, it waits for 10 seconds before it resends that given LSA.

On R2

R2(config)#int f0/0
R2(config-if)#**ip ospf retransmit-interval** 10

When an OSPF enabled router sends an LSA to it's neighbor, it keeps the LSA until it receives an ACK from that given neighbor. If the retransmission timer expires and the router receives no ACKs, the router will resend that LSA. The default timer is set to 5 seconds, and the range is 1 – 65535.

Task 6

Configure R2 such that it limits the number of non-self generated LSAs that an OSPF routing process can keep in the OSPF LSDB to 900.

On R2

R2(config-if)#router ospf 1
R2(config-router)#**max-lsa** 900

By default the number of non-self-generated LSAs that an OSPF routing process can keep in the database is not limited. To limit this number, we can use the "max-lsa" command in the router configuration mode. This command contains the following fields:

Max-lsa maximum-number [threshold-percentage] [warning-only] [ignore-time minutes] [ignore-count #] [reset-time minutes]

- **Maximum-number** – The maximum number of non-self-generated LSAs that an OSPF router can keep in the OSPF database.
- **Threshold-percentage** – The percentage of the maximum LSA number, as specified by the maximum-number, at which a warning message is logged. The default is 75.
- **Warning-only** – This specifies that only a warning message is sent when the maximum limit for LSAs is exceeded.
- **Ignore-time** ***minutes*** – This value specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes.
- **Ignore-count count-number** – Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times. The router can not exceed this number.
- **Reset-time** ***minutes*** - This value specifies the time, in minutes, after which the ignore count is reset to zero. The default is 10 minutes.

Task 7

R3 and R4 should exchange hellos every 15 seconds with a dead interval of 60 seconds. Do NOT use ip ospf dead-interval to accomplish this task.

On R3

R3(config-router)#Int S0/0.34
R3(config-subif)#**ip ospf hello-interval 15**

On R4

R4(config-router)#Int S0/0.43
R4(config-subif)#**ip ospf hello-interval 15**

Once the hello-interval is set, OSPF process will set the dead-interval to be four times the hello-interval. The default value for the hello timer is as follows:

- **On Ethernet segment its set to 10 seconds.**
- **On Non-broadcast networks, its set to 30 seconds.**

To verify the configuration:

On R4

```
R4#Show ip ospf int S0/0.43

Serial0/0.43 is up, line protocol is up
  Internet Address 10.1.34.4/24, Area 0
  Process ID 1, Router ID 4.4.4.4, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 15, Dead 60, Wait 60, Retransmit 5
    oob-resync timeout 60
    Hello due in 00:00:03
  Supports Link-local Signaling (LLS)
   (The rest of the output is omitted)
```

Task 9

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 3 – OSPF Authentication

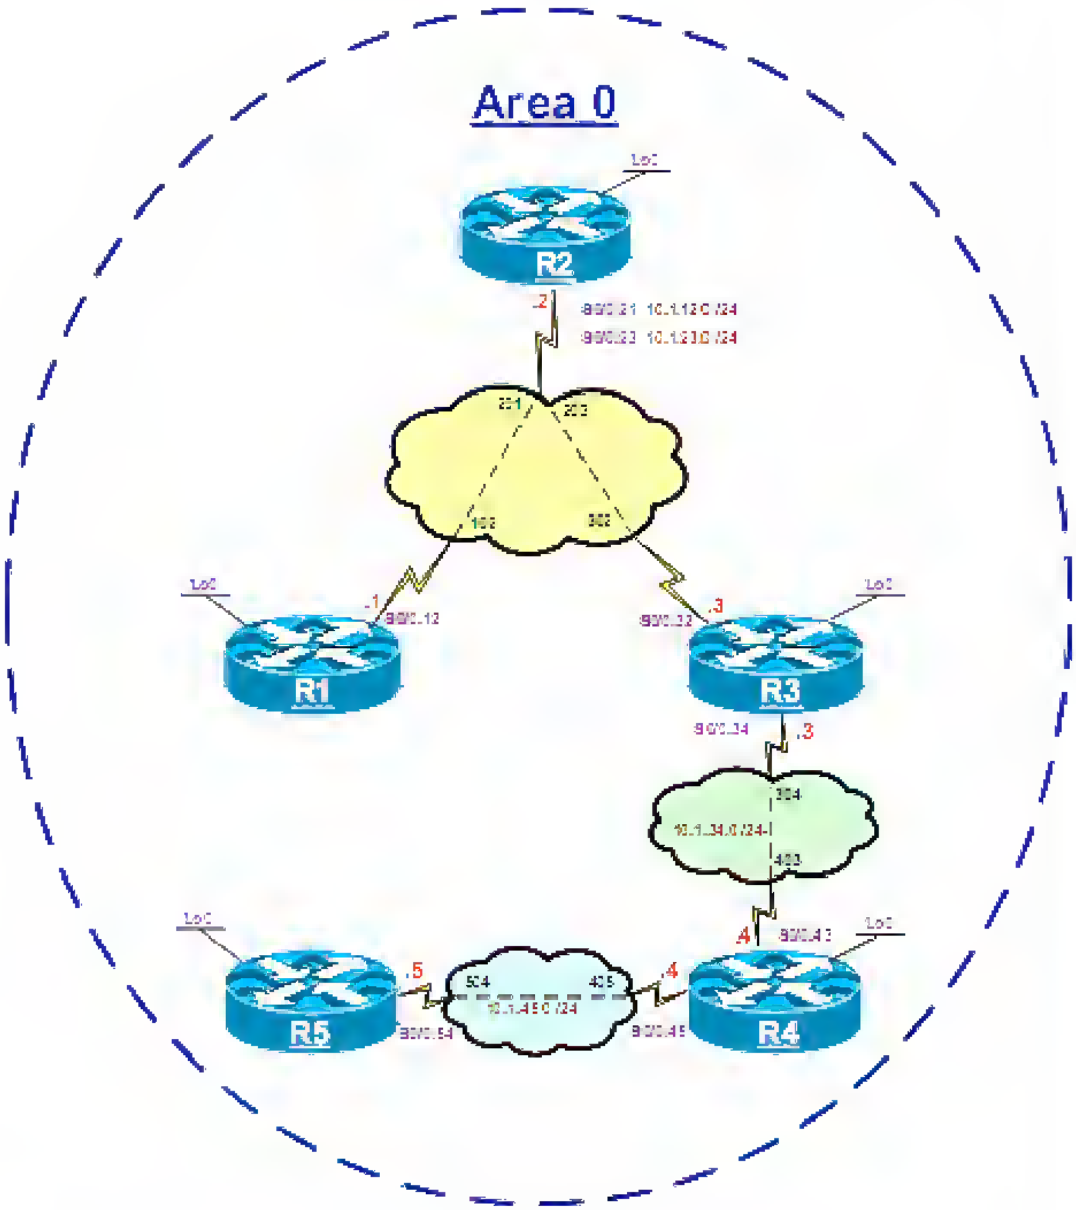

Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Use the IP addressing scheme below for IP addressing assignment.

IP Addressing scheme:

Routers	Interface / IP Address	Connecting to:
R1	S0/0.12 – 10.1.12.1 /24 Loopback 0 – 1.1.1.1 /24	R2
R2	S0/0.21 – 10.1.12.2 /24 S0/0.23 – 10.1.23.2 /24 Loopback 0 – 2.2.2.2 /24	R1 R3
R3	S0/0.32 – 10.1.23.3 /24 S0/0.34 – 10.1.34.3 /24 Loopback 0 – 3.3.3.3 /24	R2 R4
R4	S0/0.43 – 10.1.34.4 /24 S0/0.45 – 10.1.45.4 /24 Loopback 0 – 4.4.4.4 /24	R3 R5
R5	S0/0.54 – 10.1.45.5 /24 Loopback 0 – 5.5.5.5 /24	R4

Task 1

Configure the frame-relay interface/s and the loopback interface/s of all routers in area 0, and ensure that the loopback interfaces are advertised with their correct mask. The router-id of the routers in this area should be based on their loopback 0 interfaces' IP address.

On R1

```
R1(config-fr-dlci)#int lo0
R1(config-if)#ip ospf net point-to-point

R1(config)#Router ospf 1
R1(config-router)#router-id 1.1.1.1

R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 0
```

On R2

```
R2(config)#int lo0
R2(config-if)#ip ospf network point-to-point

R2(config-if)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0

On R3

R3(config)#int lo0
R3(config-if)#ip ospf network point-to-point

R3(config-if)#router ospf 1
R3(config-router)#router-id 3.3.3.3

R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 0

On R4

R4(config)#int lo0
R4(config-if)#ip ospf network point-to-point

R4(config-if)#router ospf 1
R4(config-router)#router-id 4.4.4.4

R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 0

On R5

R5(config)#int lo0
R5(config-if)#ip ospf network point-to-point

R5(config-if)#router ospf 1
R5(config-router)#router-id 5.5.5.5

R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 0

To verify the configuration:

On R1

R1#Show ip route ospf | Inc O

```
O       2.2.2.0 [110/65] via 10.1.12.2, 00:00:21, Serial0/0.12
O       3.3.3.0 [110/129] via 10.1.12.2, 00:00:21, Serial0/0.12
O       4.4.4.0 [110/193] via 10.1.12.2, 00:00:21, Serial0/0.12
O       5.5.5.0 [110/257] via 10.1.12.2, 00:00:10, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:00:21, Serial0/0.12
O       10.1.45.0 [110/256] via 10.1.12.2, 00:00:21, Serial0/0.12
O       10.1.34.0 [110/192] via 10.1.12.2, 00:00:21, Serial0/0.12
```

Task 2

Configure plain text authentication on all the Frame-relay links in this area. You should use a sub-router configuration command as part of the solution to this task. Use "Cisco" as the password for this authentication.

OSPF supports two types of authentication, plain text (64 bit password) and MD5 (Which consists of a key ID and 128 bit password). In OSPF, authentication must be enabled and then applied.
In OSPF, enabling authentication can be configured in two different ways; one way to enable OSPF authentication is to configure it in the router configuration mode, in which case authentication is enabled globally on all OSPF enabled interfaces in the specified area. The second way to enable authentication is to configure it directly on the interface for which authentication is required.

On R1

```
R1(config)#router ospf 1
R1(config-router)#area 0 authentication

R1(config-router)#int S0/0.12
R1(config-subif)#ip ospf authentication-key Cisco
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#area 0 authentication

R2(config-router)#int S0/0.21
R2(config-subif)#ip ospf authentication-key Cisco

R2(config-subif)#int S0/0.23
R2(config-subif)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R2

R2#Show ip ospf interface S0/0.21

```
Serial0/0.21 is up, line protocol is up
  Internet Address 10.1.12.2/24, Area 0
  Process ID 1, Router ID 2.2.2.2, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:06
  Supports Link-local Signaling (LLS)
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1
  Suppress hello for 0 neighbor(s)
  Simple password authentication enabled
```

Note the output of the above "Show" command verifies that a simple password authentication is enabled and applied to this interface.

R2#Show ip ospf neighbor

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.1           0   FULL/  -        00:00:30    10.1.12.1       Serial0/0.21
```

R2#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:05:00, Serial0/0.21
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#area 0 authentication

R3(config)#int S0/0.32
R3(config-subif)#ip ospf authentication-key Cisco

R3(config)#int S0/0.34
```

R3(config-subif)#**ip ospf authentication-key** Cisco

To verify the configuration:

On R3

R3#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/129] via 10.1.23.2, 00:01:36, Serial0/0.32
O       2.2.2.0 [110/65] via 10.1.23.2, 00:01:36, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:01:36, Serial0/0.32
```

On R4

R4(config)#int S0/0.43
R4(config-subif)#**ip ospf authentication-key** Cisco

R4(config-subif)#int S0/0.45
R4(config-subif)#**ip ospf authentication-key** Cisco

R4(config-subif)#router ospf 1
R4(config-router)#**area 0 authentication**

To verify the configuration:

On R4

R4#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/193] via 10.1.34.3, 00:00:21, Serial0/0.43
O       2.2.2.0 [110/129] via 10.1.34.3, 00:00:21, Serial0/0.43
O       3.3.3.0 [110/65] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.12.0 [110/192] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:21, Serial0/0.43
```

On R5

R5(config)#Router ospf 1
R5(config-router)#**area 0 authentication**

R5(config-router)#int S0/0.54
R5(config-subif)#**ip ospf authentication-key** Cisco

To verify the configuration:

On R5

R5#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/257] via 10.1.45.4, 00:00:44, Serial0/0.54
O       2.2.2.0 [110/193] via 10.1.45.4, 00:00:44, Serial0/0.54
O       3.3.3.0 [110/129] via 10.1.45.4, 00:00:44, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:00:44, Serial0/0.54
O       10.1.12.0 [110/256] via 10.1.45.4, 00:00:44, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:00:44, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:00:44, Serial0/0.54
```

Task 3

Remove the authentication configuration from the previous task and ensure that every router sees every route advertised in area 0.

On All Routers

```
(config)#router ospf 1
(config-router)#NO area 0 authentication
```

On R1

```
R1(config)#int S0/0.12
R1(config-subif)#NO ip ospf authentication-key Cisco
```

On R2

```
R2(config-subif)#int S0/0.21
R2(config-subif)#NO ip ospf authentication-key Cisco

R2(config-subif)#int S0/0.23
R2(config-subif)#NO ip ospf authentication-key Cisco
```

On R3

```
R3(config-router)#int S0/0.32
```

R3(config-subif)#NO ip ospf authentication-key Cisco

R3(config-subif)#int S0/0.34
R3(config-subif)#NO ip ospf authentication-key Cisco

On R4

R4(config)#int S0/0.43
R4(config-subif)#NO ip ospf authentication-key Cisco

R4(config-subif)#int S0/0.45
R4(config-subif)#NO ip ospf authentication-key Cisco

On R5

R5(config)#int S0/0.54
R5(config-subif)#NO ip ospf authentication-key Cisco

To verify the configuration:

On R1

R1#Show ip route ospf | Inc O

```
O       2.2.2.0 [110/65] via 10.1.12.2, 00:00:10, Serial0/0.12
O       3.3.3.0 [110/129] via 10.1.12.2, 00:00:10, Serial0/0.12
O       4.4.4.0 [110/193] via 10.1.12.2, 00:00:10, Serial0/0.12
O       5.5.5.0 [110/257] via 10.1.12.2, 00:00:10, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:00:10, Serial0/0.12
O       10.1.45.0 [110/256] via 10.1.12.2, 00:00:10, Serial0/0.12
O       10.1.34.0 [110/192] via 10.1.12.2, 00:00:10, Serial0/0.12
```

Task 4

Configure **MD5** authentication on all the Frame-relay links in this area. You should use a sub-router configuration command as part of the solution to this task. Use "Cisco" as the password for this authentication.

The following command enables MD5 authentication on the routers using the router configuration mode:

On All Routers

(config)#router ospf 1
(config-router)#area 0 authentication message-digest

On R1

R1(config-router)#int S0/0.12
R1(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

On R2

R2(config-router)#int S0/0.21
R2(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

R2(config-subif)#int S0/0.23
R2(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

To verify the configuration:

On R2

R2#Show ip ospf interface S0/0.21

```
Serial0/0.21 is up, line protocol is up
  Internet Address 10.1.12.2/24, Area 0
  Process ID 1, Router ID 2.2.2.2, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
  Supports Link-local Signaling (LLS)
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 1.1.1.1
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
    Youngest key id is 1
```

Note the output of the above "Show" command reveals that MD5 authentication is enabled and applied and the key 1 is in use.

R2#Show ip route ospf

```
     1.0.0.0/24 is subnetted, 1 subnets
O       1.1.1.0 [110/65] via 10.1.12.1, 00:01:50, Serial0/0.21
```

On R3

```
R3(config)#int S0/0.32
R3(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

R3(config-subif)#int S0/0.34
R3(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

On R3

R3#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/129] via 10.1.23.2, 00:00:11, Serial0/0.32
O       2.2.2.0 [110/65] via 10.1.23.2, 00:00:11, Serial0/0.32
O       10.1.12.0 [110/128] via 10.1.23.2, 00:00:11, Serial0/0.32
```

On R4

```
R4(config)#int S0/0.45
R4(config-subif)#ip ospf message-digest-key 1 MD5 Cisco

R4(config-subif)#int S0/0.43
R4(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

On R4

R4#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/193] via 10.1.34.3, 00:00:21, Serial0/0.43
O       2.2.2.0 [110/129] via 10.1.34.3, 00:00:21, Serial0/0.43
O       3.3.3.0 [110/65] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.12.0 [110/192] via 10.1.34.3, 00:00:21, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:21, Serial0/0.43
```

On R5

```
R5(config)#int S0/0.54
R5(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

On R5

```
R5#Show ip route ospf | Inc O

O       1.1.1.0 [110/257] via 10.1.45.4, 00:00:42, Serial0/0.54
O       2.2.2.0 [110/193] via 10.1.45.4, 00:00:42, Serial0/0.54
O       3.3.3.0 [110/129] via 10.1.45.4, 00:00:42, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:00:42, Serial0/0.54
O       10.1.12.0 [110/256] via 10.1.45.4, 00:00:42, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:00:42, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:00:42, Serial0/0.54
```

Task 5

Remove the authentication configuration from the previous task and ensure that every router sees every route advertised in area 0.

On All Routers:

```
(config)#router ospf 1
(config-router)#NO area 0 authentication message-digest
```

On R1

```
R1(config-router)#int S0/0.12
R1(config-if)#NO ip ospf message-digest-key 1 MD5 Cisco
```

On R2

```
R2(config)#int S0/0.21
R2(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

R2(config-subif)#int S0/0.23
```

R2(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

On R3

R3(config)#int S0/0.32
R3(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

R3(config-subif)#int S0/0.34
R3(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

On R4

R4(config)#int S0/0.43
R4(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

R4(config-subif)#int S0/0.45
R4(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

On R5

R5(config)#int S0/0.54
R5(config-subif)#NO ip ospf message-digest-key 1 MD5 Cisco

To verify the configuration:

On R5

R5#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/257] via 10.1.45.4, 00:09:13, Serial0/0.54
O       2.2.2.0 [110/193] via 10.1.45.4, 00:09:13, Serial0/0.54
O       3.3.3.0 [110/129] via 10.1.45.4, 00:09:13, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:09:13, Serial0/0.54
O       10.1.12.0 [110/256] via 10.1.45.4, 00:09:13, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:09:13, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:09:13, Serial0/0.54
```

R5#Show run | S router ospf 1

```
router ospf 1
 router-id 5.5.5.5
 log-adjacency-changes
 network 5.5.5.5 0.0.0.0 area 0
 network 10.1.45.5 0.0.0.0 area 0
```

Task 6

Configure MD5 authentication on the Frame-relay link connecting R1 to R2, you should use a router configuration command as part of the solution to this task. The password should be "ccie".

On Both Routers:

(config)#router ospf 1
(config-router)#area 0 authentication message-digest

On R1

R1(config)#int S0/0.12
R1(config-subif)#ip ospf message-digest-key 1 MD5 ccie

On R2

R2(config)#int S0/0.21
R2(config-subif)#ip ospf message-digest-key 1 MD5 ccie

To verify the configuration:

On R2

R2#Show ip route ospf | Inc O

O 1.1.1.0 [110/65] via 10.1.12.1, 00:00:43, Serial0/0.21

Note because authentication was enabled in router configuration mode, every router in area 0 MUST have authentication enabled. Since R3 does NOT have authentication enabled, R2 will NOT form an adjacency with R3, therefore, they will NOT exchange updates.

Task 7

Configure these routers such that every router has every prefix advertised in this topology in their routing table and Link state database.

On R3, R4 and R5

(config-router)#area 0 authentication message-digest

To verify the configuration:

On R2

R2#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:01:01, Serial0/0.21
O       3.3.3.0 [110/65] via 10.1.23.3, 00:01:01, Serial0/0.23
O       4.4.4.0 [110/129] via 10.1.23.3, 00:01:01, Serial0/0.23
O       5.5.5.0 [110/193] via 10.1.23.3, 00:01:01, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:01:01, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:01:01, Serial0/0.23
```

Note once the authentication is enabled on the other routers, they will form adjacency and exchange routes.

Task 8

Remove the configuration from the previous task and reconfigure R2 such that every router has every prefix advertised in this topology in their routing table and Link state database. DO NOT remove the authentication that is applied to the link between R1 and R2.

On R3, R4 and R5

(config-router)#No area 0 authentication message-digest

To verify the configuration:

On R2

R2#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:00:06, Serial0/0.21
```

To configure this task, we must disable authentication on the interface facing R3 using the "IP OSPF authentication null" interface configuration command, meaning that there is no need to have authentication passed .23 interface of R2. Therefore,

R3, R4 and/or R5 do NOT need to have authentication enabled.

On R2

R2(config)#int S0/0.23
R2(config-subif)#**ip ospf authentication null**

To verify the configuration:

On R2

R2#Show ip route ospf | Inc O

O 1.1.1.0 [110/65] via 10.1.12.1, 00:03:02, Serial0/0.21
O 3.3.3.0 [110/65] via 10.1.23.3, 00:03:02, Serial0/0.23
O 4.4.4.0 [110/129] via 10.1.23.3, 00:03:02, Serial0/0.23
O 5.5.5.0 [110/193] via 10.1.23.3, 00:03:02, Serial0/0.23
O 10.1.45.0 [110/192] via 10.1.23.3, 00:03:02, Serial0/0.23
O 10.1.34.0 [110/128] via 10.1.23.3, 00:03:02, Serial0/0.23

Task 9

Re-configure the authentication password configured in task 6 to be "CCIERS" without interrupting the links operation.

To see the current configuration:

On R1

R1#Show run int S0/0.12 | b interface

interface Serial0/0.12 point-to-point
 ip address 10.1.12.1 255.255.255.0

 ip ospf message-digest-key 1 md5 ccie
 frame-relay interface-dlci 102

On R2

R2#Show run inter S0/0.21 | b interface

```
interface Serial0/0.21 point-to-point
 ip address 10.1.12.2 255.255.255.0
 ip ospf message-digest-key 1 md5 ccie
 frame-relay interface-dlci 201
```

R2#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:06:26, Serial0/0.21
O       3.3.3.0 [110/65] via 10.1.23.3, 00:06:26, Serial0/0.23
O       4.4.4.0 [110/129] via 10.1.23.3, 00:06:26, Serial0/0.23
O       5.5.5.0 [110/193] via 10.1.23.3, 00:06:26, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:06:26, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:06:26, Serial0/0.23
```

To change the passwords without any interruption to the link the second key is entered with the required password:

On R1

```
R1(config)#int S0/0.12
R1(config-subif)#ip ospf message-digest-key 2 MD5 CCIERS
```

To verify the configuration:

On R1

R1#Show ip ospf inter S0/0.12 | b Message

```
  Message digest authentication enabled
    Youngest key id is 2
    Rollover in progress, 1 neighbor(s) using the old key(s):
      key id 1
```

Note even though the second key (key 2) is only configured on R1, R1 and R2 are still authenticating based on the first key (key 1), this is revealed in the second line. But the router knows that the second key is configured (The second line in the above display) and it knows that the rollover is in progress (The third line), but the other end (R2) has not been configured yet.

On R2

```
R2(config-subif)#int S0/0.21
R2(config-if)#ip ospf message-digest-key 2 MD5 CCIERS
```

To verify the configuration:

On R2

R2#Sh ip ospf inter S0/0.21 | b Message

```
  Message digest authentication enabled
    Youngest key id is 2
```

Note once R2 is configured, both routers (R1 and R2) will switchover and use the second key for their authentication.

On R1

R1#Show ip ospf interface S0/0.12 | b Message

```
  Message digest authentication enabled
    Youngest key id is 2
```

Once R1 and R2 rollover is completed and both routers display the same youngest key without the "rollover in progress" message, we can safely remove the prior key, in this case key id 1. Remember that the newest key is NOT determined based on the numerical higher value.

On R1

R1#Show run int S0/0.12 | Inc ip ospf

```
 ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIERS
```

```
R1(config)#int S0/0.12
R1(config-subif)#NO ip ospf message-digest-key 1 md5 ccie
```

On R2

R2#Show run int S0/0.21 | Inc ip ospf

```
 ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIERS
```

```
R2(config)#int S0/0.21
R2(config-subif)#NO ip ospf message-digest-key 1 md5 CCIERS
```

Task 10

Remove the configuration from the previous task and task 8 and reconfigure MD5 authentication between R1 and R2 such that every router has every prefix advertised in this topology in their routing table and Link state database. DO NOT use any router configuration mode command to accomplish this task.

On R1 and R2

(config)#router ospf 1
(config-router)#NO area 0 authentication message-digest

On R2

R2(config)#int S0/0.23
R2(config-subif)#NO ip ospf authentication null

Note the following command enables authentication directly under the .21 interface and NOT in router configuration mode

R2(config)#int S0/0.21
R2(config-subif)#ip ospf authentication message-digest

On R1

R1(config)#int S0/0.12
R1(config-subif)#ip ospf authentication message-digest

To verify the configuration:

On R1

R1#Show run inter S0/0.12 | Inc ip ospf

ip ospf authentication message-digest
ip ospf message-digest-key 2 md5 CCIERS

On R2

R2#Show run int S0/0.21 | Inc ip ospf

ip ospf authentication message-digest
ip ospf message-digest-key 2 md5 CCIERS

Note when authentication is enabled directly under a given interface, it no longer needs to be enabled on all other routers in that area. When authentication is enabled directly under a given interface, it's called per-interface authentication.

<u>To test the configuration:</u>

<u>On R1</u>

<u>R1#Show ip route ospf | inc O</u>

```
O       2.2.2.0 [110/65] via 10.1.12.2, 00:14:36, Serial0/0.12
O       3.3.3.0 [110/129] via 10.1.12.2, 00:14:36, Serial0/0.12
O       4.4.4.0 [110/193] via 10.1.12.2, 00:14:36, Serial0/0.12
O       5.5.5.0 [110/257] via 10.1.12.2, 00:14:36, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:14:36, Serial0/0.12
O       10.1.45.0 [110/256] via 10.1.12.2, 00:14:36, Serial0/0.12
O       10.1.34.0 [110/192] via 10.1.12.2, 00:14:36, Serial0/0.12
```

<u>On R2</u>

<u>R2#Show ip route ospf | Inc O</u>

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:17:32, Serial0/0.21
O       3.3.3.0 [110/65] via 10.1.23.3, 00:17:32, Serial0/0.23
O       4.4.4.0 [110/129] via 10.1.23.3, 00:17:32, Serial0/0.23
O       5.5.5.0 [110/193] via 10.1.23.3, 00:17:32, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:17:32, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:17:32, Serial0/0.23
```

<u>Task 11</u>

Re-configure the routers using the following chart, Configure OSPF router-id of the routers to be based on their Loopback interfaces' IP address, ensure that every router has every prefix advertised in this routing domain in their routing table and Link state database:

Router	Interface	Area
R1	S0/0.12 Loopback 0	0 0
R2	S0/0.21 S0/0.23	0 1

	Loopback 0	1
R3	S0/0.32 S0/0.34 Loopback 0	1 2 2
R4	S0/0.43 S0/0.45 Loopback 0	2 3 3
R5	S0/0.54 Loopback 0	3 3

On R1

R1(config)#NO router ospf 1

R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1(config-router)#netw 10.1.12.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config)#int S0/0.12
R1(config-sub)#NO ip ospf message-digest key 2 CCIERS

On R2

R2(config)#NO router ospf 1

R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2

R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 1
R2(config-router)#area 1 virtual-link 3.3.3.3

R2(config)#int S0/0.21
R2(config-sub)#NO ip ospf message-digest key 2 CCIERS

On R3

R3(config)#NO router ospf 1
R3(config)#router ospf 1
R3(config-router)#router-id 3.3.3.3

R3(config-router)#netw 10.1.23.3 0.0.0.0 area 1
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 2
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 2
R3(config-router)#area 2 virtual-link 4.4.4.4
R3(config-router)#area 1 virtual-link 2.2.2.2

On R4

R4(config)#**NO** router ospf 1

R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4

R4(config-router)#netw 10.1.45.4 0.0.0.0 area 3
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 3
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 2

R4(config-router)#area 2 virtual-link 3.3.3.3

On R5

R5(config)#**NO** router ospf 1

R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5

R5(config-router)#netw 10.1.45.5 0.0.0.0 area 3
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 3

To verify the configuration:

On R1

R1#Show ip route ospf | Inc O

O IA 2.2.2.0 [110/65] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA 3.3.3.0 [110/129] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA 4.4.4.0 [110/193] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA 5.5.5.0 [110/257] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA 10.1.23.0 [110/128] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA 10.1.45.0 [110/256] via 10.1.12.2, 00:00:32, Serial0/0.12
O IA 10.1.34.0 [110/192] via 10.1.12.2, 00:00:32, Serial0/0.12

On R5

R5#Show ip route ospf | Inc O

```
O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:01:50, Serial0/0.54
O IA    2.2.2.0 [110/193] via 10.1.45.4, 00:01:50, Serial0/0.54
O IA    3.3.3.0 [110/129] via 10.1.45.4, 00:01:59, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:01:59, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:01:50, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:01:50, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:01:59, Serial0/0.54
```

Task 12

Configure MD5 authentication on the link between R1 and R2 in area 0, the password for this authentication should be set to Micronics, you should use router configuration mode to enable authentication.

On R1 and R2

```
(config)#router ospf 1
(config-router)#area 0 authentication message-digest
```

On R1

```
R1(config)#int S0/0.12
R1(config-subif)#ip ospf message-digest-key 1 md5 Micronics
```

On R2

```
R2(config)#int S0/0.21
R2(config-subif)#ip ospf message-digest-key 1 md5 Micronics
```

To verify the configuration:

On R2

R2#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/65] via 10.1.12.1, 00:02:32, Serial0/0.21
```

Note R2 does not have any other prefix in it's routing table, this is because authentication is enabled directly under the router configuration mode of R1 and R2 and NOT the other area 0 routers. Enter the following commands to enable authentication on the other area 0 routers:

On R3 and R4

```
(config)#router ospf 1
(config-router)#area 0 authentication message-digest
```

When a virtual-link is created on a given ABR that router becomes an area 0 router, therefore, routers R3 and R4 must have authentication enabled.

To verify the configuration:

On R5

R5#Show ip route ospf | Inc O

```
O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:06:30, Serial0/0.54
O IA    2.2.2.0 [110/193] via 10.1.45.4, 00:14:04, Serial0/0.54
O IA    3.3.3.0 [110/129] via 10.1.45.4, 00:14:04, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:14:04, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:06:30, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:14:04, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:14:04, Serial0/0.54
```

Task 13

Remove all authentications and configure MD5 authentication on the link between R1 and R2 using "Micronics" as the password. Ensure that every router in this routing domain has all the prefixes advertised by all the other routers in their routing table and link state database. You should NOT configure the other routers to accomplish this task..

On R1, R2, R3 and R4

```
(config)#router ospf 1
(config-router)#NO area 0 authentication message-digest
```

On R1

```
R1(config-router)#int S0/0.12
R1(config-subif)#ip ospf authentication message-digest
R1(config-subif)#ip ospf message-digest-key 1 Micronics
```

On R2

```
R2(config)#int S0/0.21
R2(config-subif)#ip ospf authentication message-digest
R2(config-subif)#ip ospf message-digest-key 1 Micronics
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc O

O IA    2.2.2.0 [110/65] via 10.1.12.2, 00:01:41, Serial0/0.12
O IA    3.3.3.0 [110/129] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    4.4.4.0 [110/193] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    5.5.5.0 [110/257] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    10.1.23.0 [110/128] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:01:41, Serial0/0.12
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:01:41, Serial0/0.12
```

On R5

```
R5#Show ip route ospf | Inc O

O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    2.2.2.0 [110/193] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    3.3.3.0 [110/129] via 10.1.45.4, 00:02:45, Serial0/0.54
O       4.4.4.0 [110/65] via 10.1.45.4, 00:02:45, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:02:36, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:02:45, Serial0/0.54
```

Note when configuring per-interface authentication, the other routers on the OSPF routing domain do not need to have authentication enabled.

Task 14

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 4
OSPF Cost

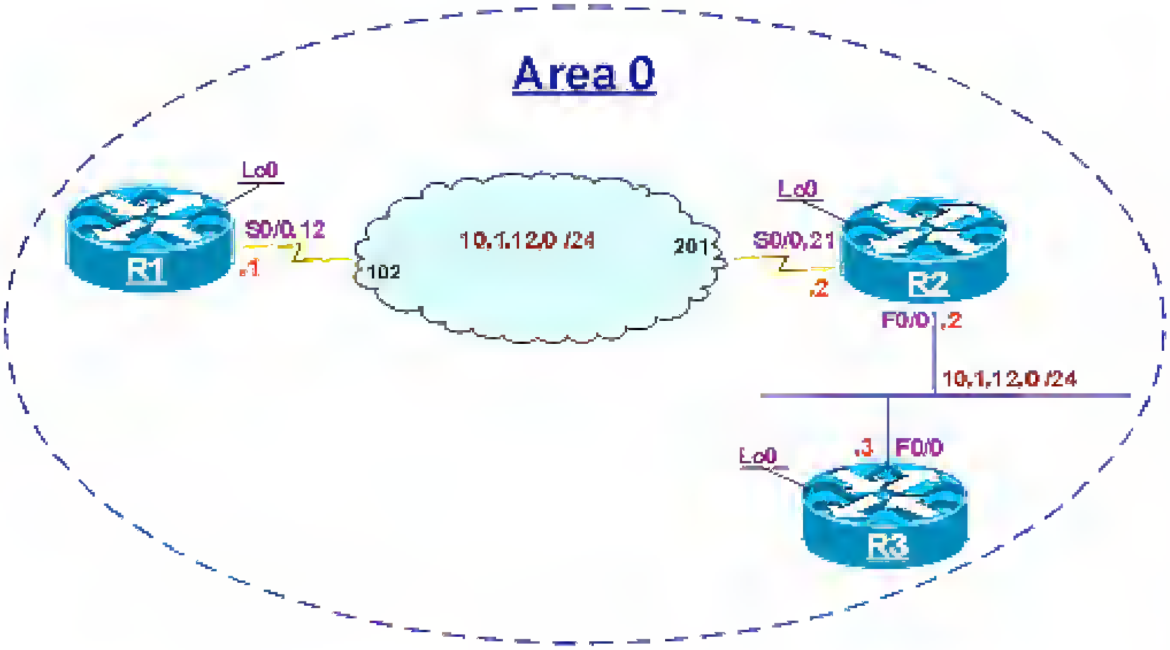

Lab Setup:

- Configure R2 and R3 in Vlan 23
- Configure the frame-relay connection between R1 and R2 in a point-to-point manner.

Ip addressing

Router	Interface and IP address
R1	Lo0 = 1.1.1.1 /8 S0/0.12 = 10.1.12.1 /24
R2	Lo0 = 2.2.2.2 /8 S0/0.21 = 10.1.12.2 /24 F0/0 = 10.1.23.2 /24
R3	Lo0 = 3.3.3.3 /8 F0/0 = 10.1.23.3 /24

Task 1

Configure all three routers in OSPF area 0 and advertise their directly connected networks in this area. Ensure that all routers have NLRI to every advertised network. Ensure that loopback interface/s is advertised with their correct mask.

On R1

R1(config)#router ospf 1
R1(config-router)#netw 0.0.0.0 0.0.0.0 are 0

R1(config-router)#int lo0
R1(config-if)#ip ospf net point-to-point

This task is asking us to ensure that the loopback interfaces are advertised with their correct mask, one way to accomplish this task is to change their network type to point-to-point.

On R2

R2(config)#router ospf 1
R2(config-router)#netw 0.0.0.0 0.0.0.0 are 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point

On R3

R3(config)#router ospf 1
R3(config-router)#netw 0.0.0.0 0.0.0.0 are 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point

To verify the configuration:

On R1

R1#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:00:16, Serial0/0.12
O    3.0.0.0/8 [110/66] via 10.1.12.2, 00:00:16, Serial0/0.12
     10.1.0.0/24 is subnetted, 2 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
O       10.1.23.0 [110/65] via 10.1.12.2, 00:00:16, Serial0/0.12
```

Task 2

Configure R1 such that it advertises a cost of 20 for it's loopback 0 interface.

You should check the cost of network 1.0.0.0 /8 that is advertised to R2 by R1.

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

O    1.0.0.0/8 [110/65] via 10.1.12.1, 00:01:24, Serial0/0.21
C    2.0.0.0/8 is directly connected, Loopback0
O    3.0.0.0/8 [110/2] via 10.1.23.3, 00:01:24, FastEthernet0/0
     10.1.0.0/24 is subnetted, 2 subnets
C       10.1.12.0 is directly connected, Serial0/0.21
C       10.1.23.0 is directly connected, FastEthernet0/0
```

Note, the cost of the loopback interface is 65; this is the result of adding the cost of the serial interface (100,000,000 / 1,544,000 = 64, remember to drop the decimal points) plus the cost of the loopback interface (100,000,000 / 8000,000,000 = 1,

remember that you can't use decimals, therefore, you should round up to 1).

Enter the following to change the cost of the lo0 interface on R1:

On R1

```
R1(config-if)#int lo0
R1(config-if)#ip ospf cost 20
```

To verify the configuration:

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

O    1.0.0.0/8 [110/84] via 10.1.12.1, 00:00:07, Serial0/0.21
C    2.0.0.0/8 is directly connected, Loopback0
O    3.0.0.0/8 [110/2] via 10.1.23.3, 00:00:07, FastEthernet0/0
     10.1.0.0/24 is subnetted, 2 subnets
C       10.1.12.0 is directly connected, Serial0/0.21
C       10.1.23.0 is directly connected, FastEthernet0/0
```

Note, the cost after the configuration is 84, which is the sum of 64 (The cost of the serial interface) plus 20 (Which is the cost of the lo0 interface).

Task 3

In the future a gigabit interface will be installed on one of the routers in this routing domain. Ensure that the costs of the other interfaces are adjusted proportionally.

On All Routers

```
(config-if)#router ospf 1
(config-router)#auto-cost reference-bandwidth 1000

#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

By default, OSPF calculates the cost of an interface by dividing the bandwidth of the interface into 100 million. Using the default value, when your network has interfaces with a bandwidth greater than 100 million is not recommended, because, OSPF will not be able to differentiate between 100mbps interface and an interface with a bandwidth that is greater than 100mbps, "IP OSPF COST" command enables you to change the OSPF cost for an interface, but a better way to accomplish this is to change the default reference value used to calculate the OSPF cost of an interface. This value can be modified using the command "auto-cost reference-bandwidth". If you are planning to use this command, remember that every router in the OSPF routing domain must be configured as well.

Task 4

Remove the command configured in task3.

On All Routers

```
(config-if)#router ospf 1
(config-router)#NO auto-cost reference-bandwidth 1000
(config)#End

#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

Task 5

Configure the routers such that the Fast Ethernet interface of these routers will have a cost of 70. The other interfaces should have their cost calculated proportionally.

On All Routers

(config-router)#router ospf 1
(config-router)#**auto-cost reference-bandwidth 7000**

The equation used by OSPF is as follows:

Reference/Bandwidth = Cost

Rearranging the formula, we get the following equation:

Reference = Cost X Bandwidth = 70 X 100,000,000 = 7,000,000,000

Reference is in units of Mbps, and by default its set to 100 which means 100,000,000 bps, now we have to divide the result by 1000,000 to get the actual reference, which is 7000.

So the cost reference should be changed to 7000

To verify the configuration:

On R3

R3#Show ip ospf int f0/0

```
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.1.23.3/24, Area 0
  Process ID 1, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 70
  Transmit Delay is 1 sec, State BDR, Priority 1
```
(The rest of the output is omitted)

Task 6

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 5
OSPF Summarization

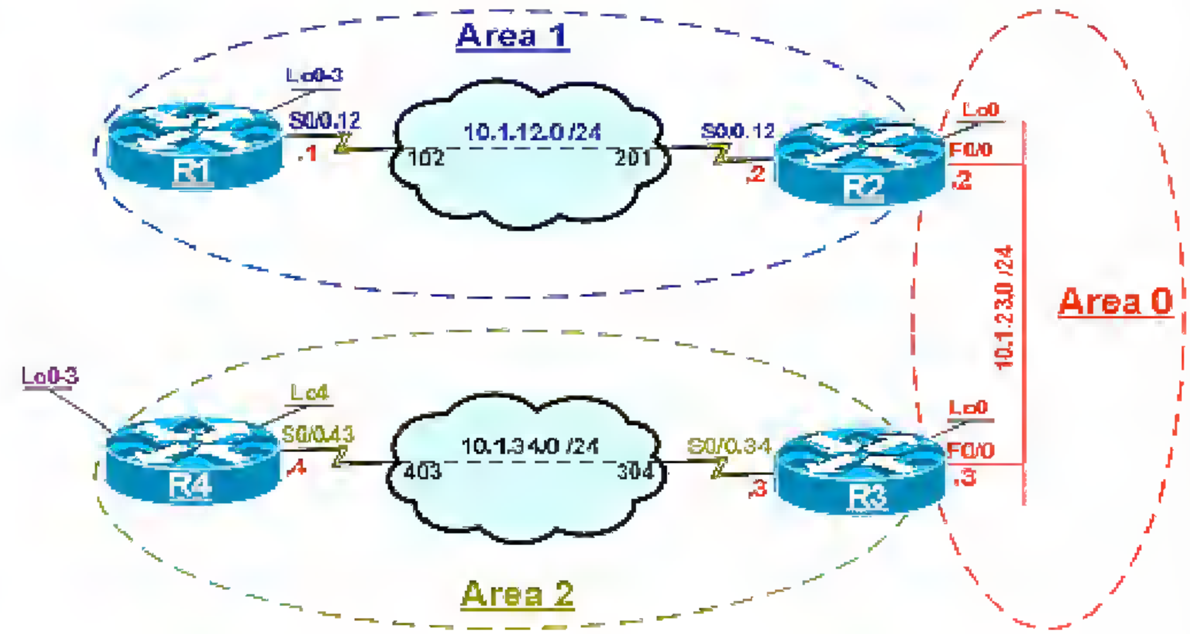

Lab setup:

- Configure R2' F0/0 and R3's F0/0 in VLAN 23
- Configure the frame-relay connection between R1, R2 and R3, R4 in a point-to-point manner.

IP Addressing

Router	Interface	IP address	Area
R1	Lo0	1.1.0.1 /24	Area 1
	Lo1	1.1.1.1 /24	Area 1
	Lo2	1.1.2.1 /24	Area 1
	Lo3	1.1.3.1 /24	Area 1
	S0/0.12	10.1.12.1 /24	Area 1
R2	Lo0	2.2.2.2 /8	Area 0
	S0/0.21	10.1.12.2 /24	Area 1
	F0/0	10.1.23.2 /24	Area 0
R3	Lo0	3.3.3.3 /8	Area 0
	S0/0.34	10.1.34.3 /24	Area 2
	F0/0	10.1.23.3 /24	Area 0
R4	Lo0	4.4.0.4 /24	External
	Lo1	4.4.1.4 /24	External
	Lo2	4.4.2.4 /24	External
	Lo3	4.4.3.4 /24	External
	Lo4	4.4.4.4 /24	Area 2
	S0/0.43	10.1.34.4 /24	Area 2

Task 1

Configure the routers as follows:

- R4 should redistribute the four loopback interfaces (4.4.0.4 /24 – 4.4.3.4 /24) in the OSPF routing domain.
- R4 should advertise it's Loopback 4 and Frame-relay interface in Area 2.
- R1 should advertise all of its interfaces in OSPF area 1.
- R2 should advertise its Loopback0, F0/0 interface in area 0 and the frame-relay interface in area 1.
- R3 should advertise its Loopback0, F0/0 interface in area 0, and its frame-relay interface in area 2.

On R1

```
R1(config)#router ospf 1
R1(config-router)#netw 0.0.0.0 0.0.0.0 area 1
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 0
```

R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 1

On R3

R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 2

On R4

R4(config)#access-list 4 permit 4.4.0.0 0.0.3.255

R4(config)#route-map TEST permit 10
R4(config-route-map)#match ip addr 4

R4(config-if)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 10.1.34.4 0.0.0.0 are 2
R4(config-router)#redistribute connected subnets route-map TEST

When redistributing routes into OSPF, the subnets keyword will redistribute all the subnets into OSPF, if this keyword is omitted, then only classful networks will be redistributed into OSPF.

Task 2

Configure the OSPF routers such that the external routes are summarized.

On R4

R4(config)#router ospf 1
R4(config-router)#summary-address 4.4.0.0 255.255.252.0

In OSPF, summarization can be configured on two types of routers: ABR/s and/or ASBRs. The internal OSPF routes can only be summarized on ABRs whereas the external (redistributed) routes can only be summarized on ASBRs. When summarizing internal routes on ABRs the "area xx range" command must be used, where xx is the area id. Summarization on ASBR can be accomplished by using the "summary-address" command.

To verify the configuration:

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 4 subnets
O IA    1.1.1.1 [110/66] via 10.1.23.2, 00:13:18, FastEthernet0/0
O IA    1.1.0.1 [110/66] via 10.1.23.2, 00:13:18, FastEthernet0/0
O IA    1.1.3.1 [110/66] via 10.1.23.2, 00:13:18, FastEthernet0/0
O IA    1.1.2.1 [110/66] via 10.1.23.2, 00:13:18, FastEthernet0/0
        2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 10.1.23.2, 00:13:50, FastEthernet0/0
C       3.0.0.0/8 is directly connected, Loopback0
        4.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O       4.4.4.4/32 [110/65] via 10.1.34.4, 00:13:19, Serial0/0.34
O E2    4.4.0.0/22 [110/20] via 10.1.34.4, 00:00:06, Serial0/0.34
        10.1.0.0/24 is subnetted, 3 subnets
O IA    10.1.12.0 [110/65] via 10.1.23.2, 00:13:19, FastEthernet0/0
C       10.1.23.0 is directly connected, FastEthernet0/0
C       10.1.34.0 is directly connected, Serial0/0.34
```

Note the external routes are summarized.

Task 3

Configure Area 1 such that networks (1.1.0.0 /24, 1.1.1.0 /24, 1.1.2.0 /24 and 1.1.3.0 /24) are summarized into the OSPF routing domain.

On R2

```
R2(config-router)#router ospf 1
R2(config-router)#area 1 range 1.1.0.0 255.255.252.0
```

Note these routes are originated by R1, but they can only be summarized by the ABR, in this topology the ABR is R2. Since the routes that are being summarized originated in area 1, the "area range" command must specify the area "area 1 range" followed by the summary network address (1.1.0.0) and then the subnet mask (255.255.252.0).

Task 4

The routers should NOT install a null 0 route in the routing table when they summarize internal or external routes.

In OSPF, the discard route is created automatically whenever a summary route is configured, there are two types of summary routes: Internal and External. When internal summary routes are configured, OSPF will inject an internal discard route, and when an external summary route is configured, the OSPF process will create an external discard route. The discard routes are created to stop forwarding loops.

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set                    This is the internal discard route

     1.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       1.1.1.1/32 [110/65] via 10.1.12.1, 00:13:53, Serial0/0.21
O       1.1.0.0/22 is a summary, 00:13:53, Null0
O       1.1.0.1/32 [110/65] via 10.1.12.1, 00:13:53, Serial0/0.21
O       1.1.3.1/32 [110/65] via 10.1.12.1, 00:13:53, Serial0/0.21
O       1.1.2.1/32 [110/65] via 10.1.12.1, 00:13:53, Serial0/0.21
C       2.0.0.0/8 is directly connected, Loopback0
```

```
        3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/2] via 10.1.23.3, 00:13:54, FastEthernet0/0
        4.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O IA    4.4.4.4/32 [110/66] via 10.1.23.3, 00:13:54, FastEthernet0/0
O E2    4.4.0.0/22 [110/20] via 10.1.23.3, 00:13:54, FastEthernet0/0
        10.1.0.0/24 is subnetted, 3 subnets
C       10.1.12.0 is directly connected, Serial0/0.21
C       10.1.23.0 is directly connected, FastEthernet0/0
O IA    10.1.34.0 [110/65] via 10.1.23.3, 00:13:55, FastEthernet0/0
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#NO discard-route internal
```

The discard route that we are discarding is the result of summarizing the internal routes, therefore we need to specify internal.

To Verify the configuration:

On R4

```
R4(config)#router ospf 1
R4(config-router)#NO discard-route external
```

In the above command we are discarding the external discard-route that was created as a result of summarizing the external routes.

Task 5

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 6
Virtual-links and GRE Tunnels

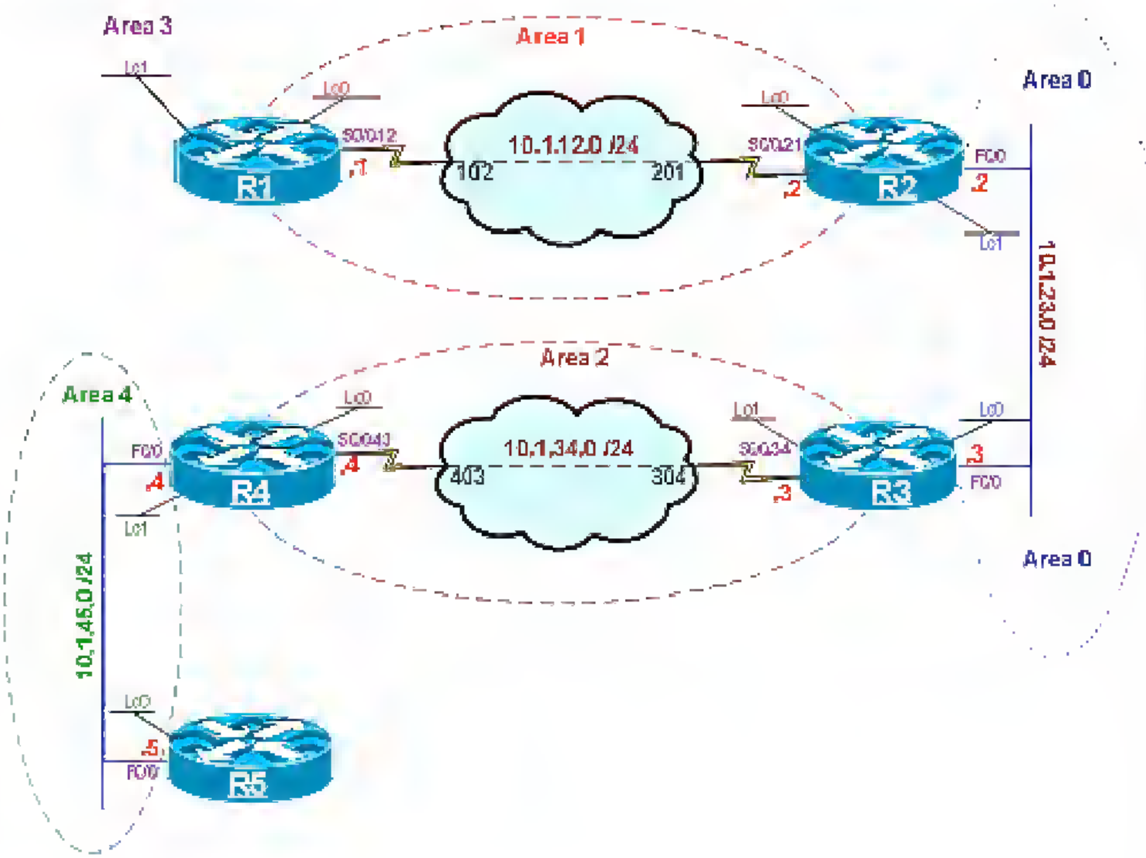

Lab Setup:

- Use the IP addressing chart below to assign IP addresses to the interfaces.
- The frame-relay connection between R1 and R2 should be configured in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- The frame-relay connection between R3 and R4 should be configured in a point-to-point manner.
- R4 and R5's F0/0 interface should be configured in VLAN 45.

IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 F/R interface	1.1.1.1 /24 10.1.1.1 / 24 10.1.12.1 /24
R2	Lo0 Lo1 F0/0 F/R interface	2.2.2.2 /24 20.2.2.2 /24 10.1.23.2 /24 10.1.12.2 /24
R3	Lo0 Lo1 F0/0 F/R interface	3.3.3.3 /24 30.3.3.3 /24 10.1.23.3 /24 10.1.34.3 /24
R4	Lo0 Lo1 F0/0 F/R interface	4.4.4.4 /24 40.4.4.4 /24 10.1.45.4 /24 10.1.34.4 /24
R5	Lo0 F0/0	5.5.5.5 /24 10.1.45.5 /24

Task 1

- R1's Loopback1 interface should be advertised in area 3 and its Frame-relay and Loopback0 interface should be advertised in area 1
- R2's Loopback0 and its frame-relay interface should be configured in area 1 and it's Loopback1 and F0/0 interface should be configured in area 0
- R3's Loopback0 and F0/0 interface should be configured in area 0 and its frame-relay and Loopback1 interface should be configured in area 2.
- R4's frame-relay and Loopback0 interface should be configured in area 2 and its F0/0 and Loopback1 interface should be configured in area 4.
- R5's Loopback and F0/0 interface should be configured in area 4.

On R1

```
R1(config-if)#router ospf 1
R1(config-router)#router-id 1.1.1.1
```

```
R1(config-router)#netw 10.1.1.1 0.0.0.0 are 3
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 1
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0
R2(config-router)#netw 20.2.2.2 0.0.0.0 are 0
```

On R3

```
R3(config-if)#router ospf 1
R3(config-router)#router-id 3.3.3.3
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 are 2
R3(config-router)#netw 30.3.3.3 0.0.0.0 are 2
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 2
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 10.1.45.4 0.0.0.0 are 4
R4(config-router)#netw 40.4.4.4 0.0.0.0 are 4
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5
R5(config-router)#netw 10.1.45.5 0.0.0.0 are 4
R5(config-router)#netw 5.5.5.5 0.0.0.0 are 4
```

Task 2

Ensure that the routes from area 3 are reachable by R1, R2, R3 and R4. Do NOT use a GRE Tunnel to accomplish this task.

Area 3 is NOT connected to area 0; the other routers won't be able to see the route advertised by this area (10.1.1.0 / 24). A virtual-link must be created that connects R1 (The ABR of area 3) to area 0.

On R1

```
R1(config)#router ospf 1
R1(config-router)#area 1 virtual-link 2.2.2.2
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#area 1 virtual-link 1.1.1.1
```

To Verify the configuration:

On R2

```
R2#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/24 is subnetted, 1 subnets
O       1.1.1.0 [110/65] via 10.1.12.1, 00:00:09, Serial0/0.21
     2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Loopback0
     3.0.0.0/24 is subnetted, 1 subnets
O       3.3.3.0 [110/2] via 10.1.23.3, 00:00:09, FastEthernet0/0
     4.0.0.0/24 is subnetted, 1 subnets
O IA    4.4.4.0 [110/66] via 10.1.23.3, 00:00:09, FastEthernet0/0
     20.0.0.0/24 is subnetted, 1 subnets
C       20.2.2.0 is directly connected, Loopback1
     10.0.0.0/24 is subnetted, 1 subnets
O IA    10.1.1.0 [110/65] via 10.1.12.1, 00:00:00, Serial0/0.21
     10.1.0.0/24 is subnetted, 3 subnets
C       10.1.12.0 is directly connected, Serial0/0.21
```

Note the route from area 3 is advertised

```
C       10.1.23.0 is directly connected, FastEthernet0/0
O IA    10.1.34.0 [110/65] via 10.1.23.3, 00:00:02, FastEthernet0/0
     30.0.0.0/24 is subnetted, 1 subnets
O IA    30.3.3.0 [110/2] via 10.1.23.3, 00:00:02, FastEthernet0/0
```

Task 3

Ensure that all the advertised networks are reachable by all the routers. Use any IP addressing and do NOT use a Virtual-link to accomplish this task.

The routing table of R5 reveals that only network 40.4.4.0 / 24 was propagated by R4. The reason for this behavior is as follows:
Area 4 does not have a connection (Logical or Physical) to area 0.

In order to rectify this problem we must create a virtual-link, since virtual-link is not allowed in this task a GRE tunnel must be used.
To display the routing table of R5 before creating a GRE tunnel:

On R5

R5#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     5.0.0.0/24 is subnetted, 1 subnets
C       5.5.5.0 is directly connected, Loopback0
     40.0.0.0/24 is subnetted, 1 subnets
O       40.4.4.0 [110/2] via 10.1.45.4, 00:04:00, FastEthernet0/0
     10.1.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, FastEthernet0/0
```

To fix this problem we must create a GRE tunnel as follows:

On R4

```
R4(config-if)#router ospf 1
R4(config-router)#netw 200.1.34.4 0.0.0.0 are 0

R4(config)#int tu1
R4(config-if)#ip addr 200.1.34.4 255.255.255.0
R4(config-if)#tunnel source 10.1.34.4
R4(config-if)#tunnel destination 10.1.34.3
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 200.1.34.3 0.0.0.0 are 0

R3(config)#int tu1
R3(config-if)#ip addr 200.1.34.3 255.255.255.0
R3(config-if)#tunnel source 10.1.34.3
R3(config-if)#tunnel destination 10.1.34.4
```

To Verify the configuration:

On R5

R5#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/24 is subnetted, 1 subnets
O IA    1.1.1.0 [110/11178] via 10.1.45.4, 00:00:07, FastEthernet0/0
     2.0.0.0/24 is subnetted, 1 subnets
O IA    2.2.2.0 [110/11114] via 10.1.45.4, 00:00:07, FastEthernet0/0
     3.0.0.0/24 is subnetted, 1 subnets
O IA    3.3.3.0 [110/11113] via 10.1.45.4, 00:00:07, FastEthernet0/0
     4.0.0.0/24 is subnetted, 1 subnets
O IA    4.4.4.0 [110/2] via 10.1.45.4, 00:00:47, FastEthernet0/0
```

```
O IA 200.1.34.0/24 [110/11112] via 10.1.45.4, 00:00:47, FastEthernet0/0
     20.0.0.0/24 is subnetted, 1 subnets
O IA    20.2.2.0 [110/11114] via 10.1.45.4, 00:00:08, FastEthernet0/0
     5.0.0.0/24 is subnetted, 1 subnets
C       5.5.5.0 is directly connected, Loopback0
     40.0.0.0/24 is subnetted, 1 subnets
O       40.4.4.0 [110/2] via 10.1.45.4, 00:00:49, FastEthernet0/0
     10.0.0.0/24 is subnetted, 1 subnets
O IA    10.1.1.0 [110/11178] via 10.1.45.4, 00:00:09, FastEthernet0/0
     10.1.0.0/24 is subnetted, 4 subnets
O IA    10.1.12.0 [110/11177] via 10.1.45.4, 00:00:09, FastEthernet0/0
O IA    10.1.23.0 [110/11113] via 10.1.45.4, 00:00:09, FastEthernet0/0
O IA    10.1.34.0 [110/65] via 10.1.45.4, 00:00:49, FastEthernet0/0
C       10.1.45.0 is directly connected, FastEthernet0/0
     30.0.0.0/24 is subnetted, 1 subnets
O IA    30.3.3.0 [110/66] via 10.1.45.4, 00:00:49, FastEthernet0/0
```

Note all the routes are advertised. The IP address of the tunnel interface MUST be advertised in area 0 or else the tunnel will not work.

Task 4

Remove the configuration from the previous task and replace it with virtual-link.

On R4

```
R4(config-if)#router ospf 1
R4(config-router)#NO netw 200.1.34.4 0.0.0.0 are 0

R4(config)#NO int tu1

R4(config)#router ospf 1
R4(config-router)#area 2 virtual-link 3.3.3.3
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#NO netw 200.1.34.3 0.0.0.0 are 0

R3(config)#NO int tu1
```

```
R3(config)#router ospf 1
R3(config-router)#area 2 virtual-link 4.4.4.4
```

Task 5

Configure a simple clear text authentication for the virtual-link that connects area 3 to area 0. Use "**Cisco**" as the password.

On R1

```
R1(config)#router ospf 1
R1(config-router)#area 1 virtual-link 2.2.2.2 authentication
R1(config-router)#area 1 virtual-link 2.2.2.2 authentication-key Cisco
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#area 1 virtual-link 1.1.1.1 authentication
R2(config-router)#area 1 virtual-link 1.1.1.1 authentication-key Cisco
```

Task 6

Configure MD5 authentication for the virtual-link that connects area 4 to area 0, use "**cisco**" as the password.

On R4

```
R4(config)#router ospf 1
R4(config-router)#area 2 virtual-link 3.3.3.3 authentication message-digest
R4(config-router)#area 2 virtual-link 3.3.3.3 message-digest-key 1 md5 cisco
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#area 2 virtual-link 4.4.4.4 authentication message-digest
R3(config-router)#area 2 virtual-link 4.4.4.4 message-digest-key 1 md5 cisco
```

Task 7

Change the password from "cisco" to "CCIE" for the virtual-link that connects area 4 to area 0, without interrupting the link.

On R4

R4(config)#router ospf 1
R4(config-router)#area 2 virtual-link 3.3.3.3 message-digest-key 2 md5 CCIE

On R3

R3(config)#router ospf 1
R3(config-router)#area 2 virtual-link 4.4.4.4 message-digest-key 2 md5 CCIE

Task 8

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 7
OSPF Stub, Totally Stubby, and NSSA Areas

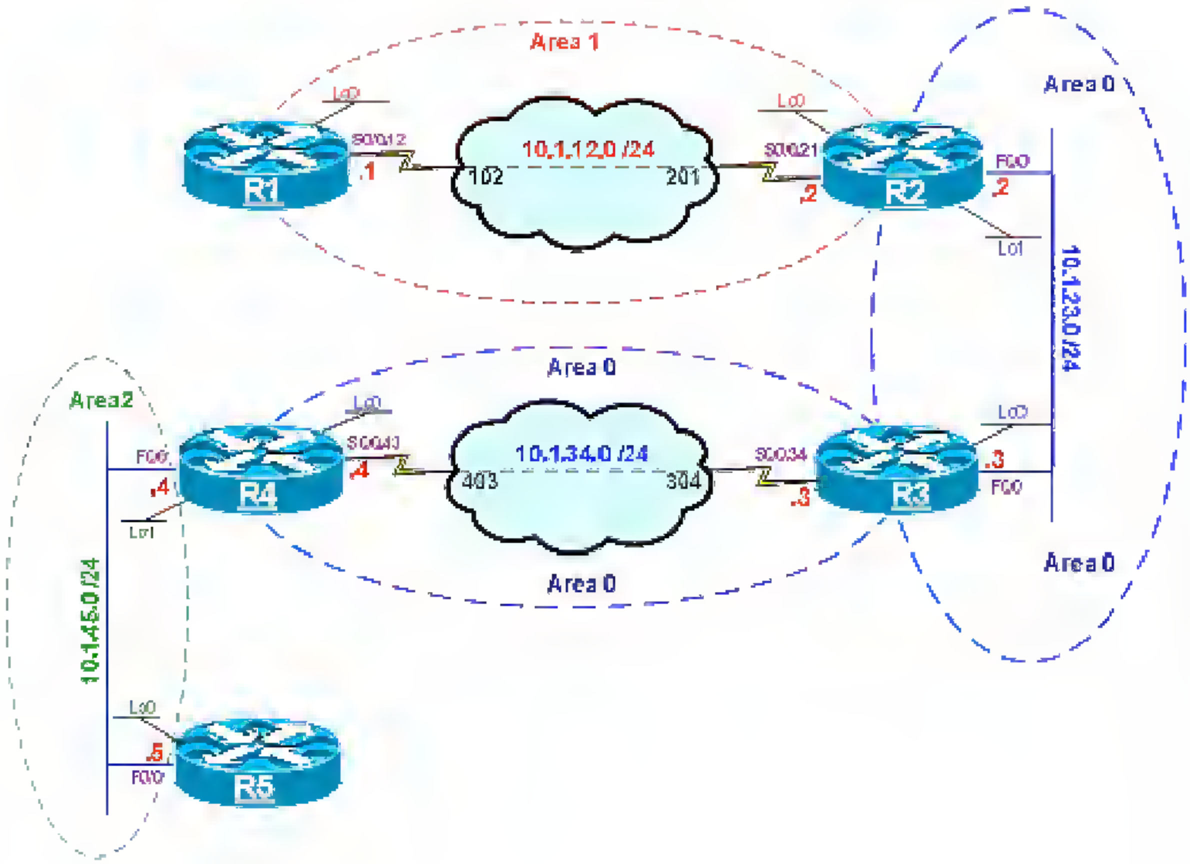

Lab Setup:

- R4 and R5's F0/0 interface should be configured in VLAN 45.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- The frame-relay connection between R1, R2 and R3, R4 should be configured in a point-to-point manner.
- Use the IP addressing chart below to assign IP addresses to the routers.

IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12	1.1.1.1 /24 10.1.12.1 /24
R2	Lo0 Lo1 F0/0 S0/0.21	2.2.2.2 /24 22.2.2.2 /24 10.1.23.2 /24 10.1.12.2 /24
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /24 10.1.23.3 /24 10.1.34.3 /24
R4	Lo0 Lo1 F0/0 S0/0.43	4.4.4.4 /24 44.4.4.4 /24 10.1.45.4 /24 10.1.34.4 /24
R5	Lo0 F0/0	5.5.5.5 /24 10.1.45.5 /24

Task 1

Configure OSPF as follows:

- Configure R1's Loopback0 and Frame-relay interface in area 1
- Configure R2's Loopback0 and Frame-relay interface in area 1 and R2's Loopback1 and F0/0 interface should be configured in area 0.
- Configure R3's Loopback0, F0/0, and Frame-relay interface in OSPF area 0.
- Configure R4's Loopback0, and Frame-relay interface in area 0, and it's Loopback1 and F0/0 in OSPF area 2.
- Configure R5's Loopback0 and F0/0 interface in OSPF area 2.
- The loopback interfaces must be advertised with their correct mask.

On R1

```
R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 1
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 1
```

```
R1(config)#Interface Lo0
R1(config-if)#Ip ospf network point-to-point
```

On R2

```
R2(config-if)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 1
R2(config-router)#netw 22.2.2.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0

R2(config)#Interface Lo0
R2(config-if)#Ip ospf network point-to-point

R2(config)#Interface Lo1
R2(config-if)#Ip ospf network point-to-point
```

On R3

```
R3(config-if)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 are 0

R3(config)#Interface Lo0
R3(config-if)#Ip ospf network point-to-point
```

On R4

```
R4(config-if)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 0
R4(config-router)#netw 10.1.34.4 0.0.0.0 are 0
R4(config-router)#netw 44.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 10.1.45.4 0.0.0.0 are 2

R4(config)#Interface Lo0
R4(config-if)#Ip ospf network point-to-point

R4(config)#Interface Lo1
R4(config-if)#Ip ospf network point-to-point
```

On R5

```
R5(config-if)#router ospf 1
```

```
R5(config-router)#netw 5.5.5.5 0.0.0.0 are 2
R5(config-router)#netw 10.1.45.5 0.0.0.0 are 2

R5(config)#Interface Lo0
R5(config-if)#Ip ospf network point-to-point
```

Task 2

Configure area 1 such that it does not receive LSA types 4 and 5.

On R1

```
R1(config)#router ospf 1
R1(config-router)#area 1 stub
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#area 1 stub
```

Important points to understand about a STUB area:

- **A STUB area can not be a transit area for Virtual link, but a GRE tunnel can be used instead.**
- **A STUB area can not have an ASBR.**
- **The back bone area can not be configured as a STUB area.**
- **Every router and the ABR of that area should have "area xx stub" command.**
- **No LSA type 5 (E1, or E2) is allowed in a STUB area, but the routers in the STUB area can connect to the External routes via the default route that is injected in the area by the ABR.**
- **By default, the cost of the default route is 1; this can be verified by "Show ip ospf", and Show ip route. The cost of the default route can be changed by "area xx default-cost cc", where xx is the area number, and cc is the desired cost**

Task 3

Area 2 should not receive LSA types 3, 4 and 5.

On R5

```
R5(config)#router ospf 1
R5(config-router)#area 2 stub
```

The above command must be configured on all the routers within this area.

On R4

```
R4(config)#router ospf 1
R4(config-router)#area 2 stub no-summary
```

The above command must only be configured on the ABR of this area.

Note you can reduce the routing table further by configuring an area as totally stubby. Since all the IA and E (Inter-area and External) routes are reached through the ABR and the ABR has injected a default route into the area, there is not reason to maintain the IA routes and they should be filtered.

Task 4

Create/configure the following loopback interfaces on R1 and redistribute them into OSPF routing domain:
Lo1 = 11.1.0.1 /24, Lo2= 11.1.1.1 /24, Lo3=11.1.2.1 /24 and Lo4 = 11.1.3.1/24

After the redistribution, area 1 should only receive and propagate LSA types 1, 2, 3 and 7. This area should not have the ability to connect to any external routes redistributed else where within this routing domain.

On R1

```
R1(config)#router ospf 1
R1(config-router)#NO area 1 stub
R1(config-router)#area 1 nssa
```

On R2

R2(config)#router ospf 1
R2(config-router)#**NO area 1 stub**
R2(config-router)#area 1 nssa

On R1

R1(config-if)#int lo1
R1(config-if)#ip addr 11.1.0.1 255.255.255.0
R1(config-if)#int lo2
R1(config-if)#ip addr 11.1.1.1 255.255.255.0
R1(config-if)#int lo3
R1(config-if)#ip addr 11.1.2.1 255.255.255.0
R1(config-if)#int lo4
R1(config-if)#ip addr 11.1.3.1 255.255.255.0

R1(config)#access-list 1 permit 11.1.0.0 0.0.3.255

R1(config)#route-map TEST permit
R1(config-route-map)#match ip addr 1

R1(config)#router ospf 1
R1(config-router)#redistribute connected route-map TEST subnets

Note when configuring an area as an NSSA area, by default the 0/0 route will not be injected by the ABR of that area.

Task 5

Create/configure the following loopback interfaces on R5 and redistribute them into OSPF routing domain:
Lo1 = 55.1.0.5 /24, Lo2= 55.1.1.5 /24, Lo3=55.1.2.5 /24 and Lo4 = 55.1.3.5 /24

After the redistribution, the routers in this area should only maintain and propagate LSA types 1, 2, 3, 7 and a default route.

On R5

R5(config)#int lo1
R5(config-if)#ip addr 55.1.0.5 255.255.255.0

R5(config-if)#int lo2

R5(config-if)#ip addr 55.1.1.5 255.255.255.0

R5(config-if)#int lo3
R5(config-if)#ip addr 55.1.2.5 255.255.255.0

R5(config-if)#int lo4
R5(config-if)#ip addr 55.1.3.5 255.255.255.0

R5(config)#access-list 1 permit 55.1.0.0 0.0.3.255

R5(config-if)#route-map TEST permit 10
R5(config-route-map)#match ip addr 1

R5(config)#router ospf 1
R5(config-router)#no area 2 stub
R5(config-router)#area 2 nssa
R5(config-router)#redistribute connected subnets route-map TEST

On R4

R4(config)#router ospf 1
R4(config-router)#**NO area 2 stub no-summary**
R4(config-router)#**area 2 nssa default-information-originate**

Note when the "area stub no-summary" command is configured and must be removed, the "no area 2 stub no-summary" command will only remove the "no-summary" part of the command. You must remember to enter the "no area 2 stub" command again to remove the entire command.

R4(config-router)#**NO are 2 stub**
R4(config-router)#**area 2 nssa default-information-originate**

Note the default-information-originate command at the end of area 2 nssa will inject a default route into the area.

Task 6

Area 1 should be changed such that it receives and propagates LSA types 1, 2, 7 plus a default route. This area should NOT maintain Inter-area routes, but must have the ability to connect to these routes.

On R2

R2(config)#router ospf 1
R2(config-router)#area 1 nssa no-summary

The "no-summary" keyword filters the summary LSAs which are the LSA type 3s.

Task 7

The default route that was injected into area 1 should have a cost of 50.

On R2

R2(config)#router ospf 1
R2(config-router)#area 1 default-cost 50

By default, the cost of the default route injected into a given area is 1; this can be verified by "Show ip route" command, remember when looking at the output of the "Show ip route" command, the cost of the default route should be 65, this is the cost of the link to the ABR (The frame-relay link) plus 1 (The default cost of the default route).
The default cost of the injected default route can be changed using the "Area xx default-cost cc", where cc is the new cost replacing the default value.
Note the new cost of the default route after configuring this task should be 114 (64 + 50).

Task 8

Erase the startup config and reload the routers before proceeding to the next lab

Lab 8 – OSPF Filtering

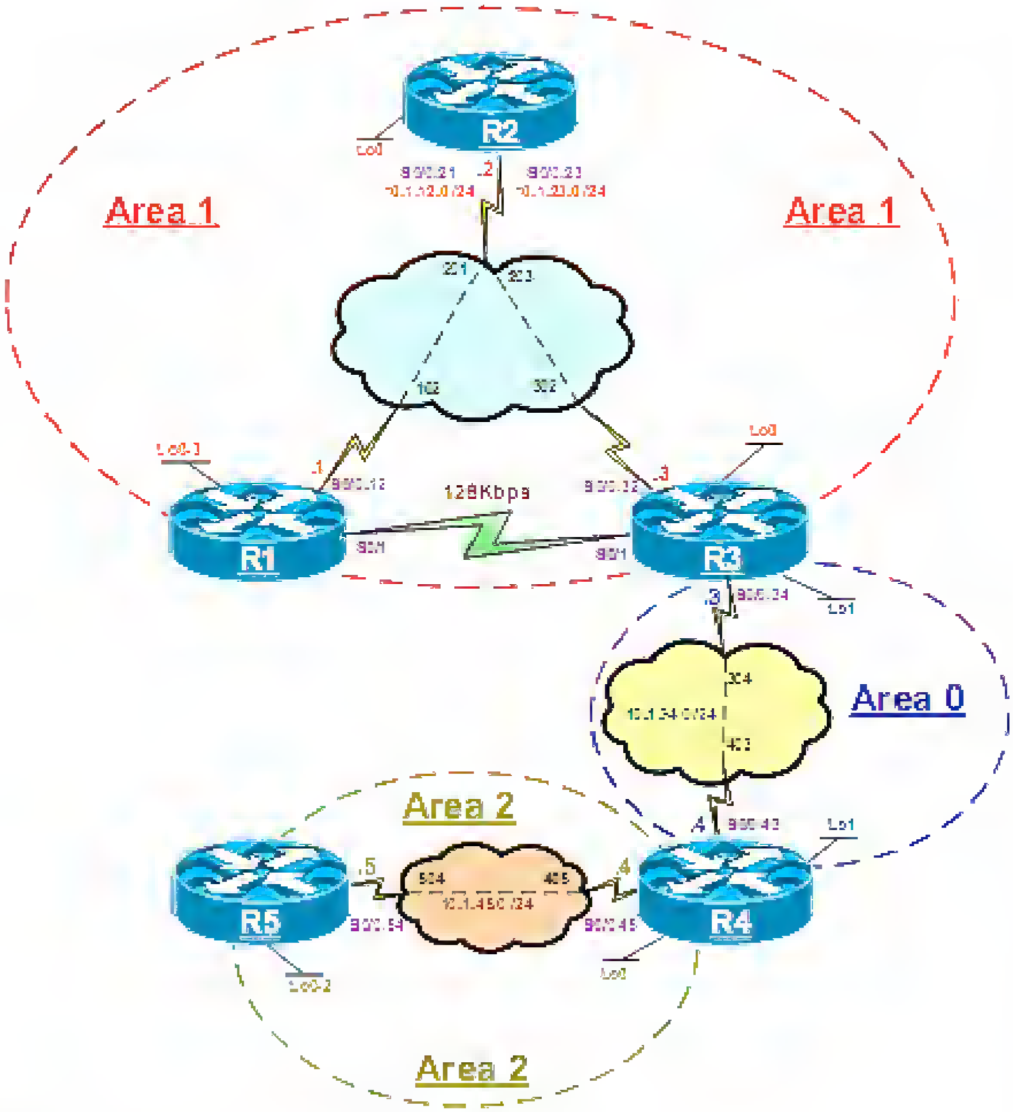

Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Configure the serial interface connecting R1 to R3 as HDLC.
- Use the IP addressing scheme below for IP addressing assignment.

IP Addressing scheme:

Routers	Interface / IP Address	Connecting to:
R1	S0/0.12 – 10.1.12.1 /24 S0/1 – 10.1.13.1 /24	R2 R3
R2	S0/0.21 – 10.1.12.2 /24 S0/0.23 – 10.1.23.2 /24	R1 R3
R3	S0/0.32 – 10.1.23.3 /24 S0/0.34 – 10.1.34.3 /24 S0/1 – 10.1.13.3 /24	R2 R4 R1
R4	S0/0.43 – 10.1.34.4 /24 S0/0.45 – 10.1.45.4 /24	R3 R5
R5	S0/0.54 – 10.1.45.5 /24	R4

Task 1

Configure R1's Frame-relay interface to R2, R1's HDLC connection to R3, R2's Frame-relay connection to R1 and R3, R3's Frame-relay connection to R2 and R3's HDLC connection to R1 in Area 1. Configure the bandwidth of S0/1 interface on R1 and R3 to 128K using the "bandwidth" command.

On R1

```
R1(config)#Int S0/1
R1(config-if)#Bandwidth 128

R1(config)#router ospf 1
R1(config-router)#netw 10.1.12.1 0.0.0.0 area 1
R1(config-router)#netw 10.1.13.1 0.0.0.0 area 1
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.12.2 0.0.0.0 area 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 1
```

On R3

```
R3(config)#Int S0/1
R3(config-if)#Bandwidth 128

R3(config)#router ospf 1
```

```
R3(config-router)#netw 10.1.13.3 0.0.0.0 area 1
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 1
```

To verify the configuration:

On R1

```
R1#Show ip route ospf

     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.23.0 [110/128] via 10.1.12.2, 00:00:44, Serial0/0.12
```

Task 2

Configure R3's frame-relay connection to R4 and R4's Frame-relay connection to R3 in Area 0.

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 0
```

To verify the configuration:

On R4

```
R4#Show ip route ospf

     10.0.0.0/24 is subnetted, 5 subnets
O IA    10.1.13.0 [110/845] via 10.1.34.3, 00:00:35, Serial0/0.43
O IA    10.1.12.0 [110/192] via 10.1.34.3, 00:00:35, Serial0/0.43
O IA    10.1.23.0 [110/128] via 10.1.34.3, 00:00:35, Serial0/0.43
```

Task 3

Configure R4's Frame-relay connection to R5 and R5's Frame-relay connection to R4 in Area 2.

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 2
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 2
```

To verify the configuration:

On R5

```
R5#Show ip route ospf

     10.0.0.0/24 is subnetted, 5 subnets
O IA    10.1.13.0 [110/909] via 10.1.45.4, 00:00:37, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:00:37, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:00:37, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:00:37, Serial0/0.54
```

Task 4

Create the following loopback interfaces on R1 and advertise them in Area 1.

Loopback 0 – 1.1.1.1 /24
Loopback 1 – 11.1.1.1 /24
Loopback 2 – 100.1.1.1 /24
Loopback 3 – 111.1.1.1 /24

On R1

```
R1(config)#int lo0
R1(config-if)#ip addr 1.1.1.1 255.255.255.0
```

```
R1(config-if)#int lo1
R1(config-if)#ip addr 11.1.1.1 255.255.255.0

R1(config-if)#int lo2
R1(config-if)#ip addr 100.1.1.1 255.255.255.0

R1(config-if)#int lo3
R1(config-if)#ip addr 111.1.1.1 255.255.255.0

R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 1
R1(config-router)#netw 11.1.1.1 0.0.0.0 area 1
R1(config-router)#netw 100.1.1.1 0.0.0.0 area 1
R1(config-router)#netw 111.1.1.1 0.0.0.0 are 1
```

To test the configuration:

On R5

R5#Show ip route ospf | inc O IA

```
O IA    1.1.1.1 [110/257] via 10.1.45.4, 00:02:52, Serial0/0.54
O IA    100.1.1.1 [110/257] via 10.1.45.4, 00:02:31, Serial0/0.54
O IA    111.1.1.1 [110/257] via 10.1.45.4, 00:02:31, Serial0/0.54
O IA    10.1.13.0 [110/909] via 10.1.45.4, 00:09:36, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:09:36, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:09:36, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:09:36, Serial0/0.54
O IA    11.1.1.1 [110/257] via 10.1.45.4, 00:02:42, Serial0/0.54
```

Task 5

Configure the router-id of the routers based on the following:

R1 – 1.1.1.1
R2 – 2.2.2.2
R3 – 3.3.3.3
R4 – 4.4.4.4
R5 – 5.5.5.5

On R1

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1#Clear ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#router-id 2.2.2.2

R2#Clear ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#router-id 3.3.3.3

R3#Clear ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#router-id 4.4.4.4

R4#Clear ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#router-id 5.5.5.5

R5#Clear ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

Task 6

Configure R2 to filter network 1.1.1.0 /24 from its routing table.

On R2

R2#Show ip route 1.1.1.0 255.255.255.0

% Subnet not in table

Note this network is not in the routing table of R2, because the loopback interfaces were NOT advertised with their correct mask. To fix this problem, we should advertise all the loopback interfaces from the previous task with their correct mask. Sometimes this can be a problem where an unbelievable amount of time is spent looking for a prefix that does not exist.

On R1

```
R1(config)#int lo0
R1(config-if)#ip ospf network point-to-point
R1(config-if)#int lo1
R1(config-if)#ip ospf network point-to-point
R1(config-if)#int lo2
R1(config-if)#ip ospf network point-to-point
R1(config-if)#int lo3
R1(config-if)#ip ospf network point-to-point
```

On R2

R2#Show ip route 1.1.1.0 255.255.255.0

```
Routing entry for 1.1.1.0/24
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 10.1.12.1 on Serial0/0.21, 00:02:15 ago
  Routing Descriptor Blocks:
  * 10.1.12.1, from 1.1.1.1, 00:02:15 ago, via Serial0/0.21
      Route metric is 65, traffic share count is 1
```

The following solution only affects the router that it's configured on, unless the filtering is done on the ABR from area 0 into other areas, in which case it will effect all routers down stream to that Area 0.

On R2

```
R2(config)#Access-list 1 deny 1.1.1.0 0.0.0.255
R2(config)#Access-list 1 permit any

R2(config)#router ospf 1
```

R2(config-router)#**distribute-list 1 in**

To verify the configuration:

On R2

R2#Show ip route ospf

```
     100.0.0.0/24 is subnetted, 1 subnets
O       100.1.1.0 [110/65] via 10.1.12.1, 00:00:30, Serial0/0.21
     111.0.0.0/24 is subnetted, 1 subnets
O       111.1.1.0 [110/65] via 10.1.12.1, 00:00:30, Serial0/0.21
     10.0.0.0/24 is subnetted, 5 subnets
O         10.1.13.0 [110/845] via 10.1.23.3, 00:00:30, Serial0/0.23
                    [110/845] via 10.1.12.1, 00:00:30, Serial0/0.21
O IA    10.1.45.0 [110/192] via 10.1.23.3, 00:00:30, Serial0/0.23
O IA    10.1.34.0 [110/128] via 10.1.23.3, 00:00:30, Serial0/0.23
     11.0.0.0/24 is subnetted, 1 subnets
O          11.1.1.0 [110/65] via 10.1.12.1, 00:00:30, Serial0/0.21
```

Note the "distribute-list in" sub-router configuration mode command can be used when filtering any type of LSA on a given router ONLY, this command ONLY filters the prefix from the local router's routing table and NOT the database. The output of the following "Show" command reveals that R3 is learning network 1.1.1.0 /24 through R2, even though this prefix is NOT in R2's routing table.

On R3

R3#Show ip route ospf

```
     1.0.0.0/24 is subnetted, 1 subnets
O       1.1.1.0 [110/129] via 10.1.23.2, 00:03:11, Serial0/0.32
     100.0.0.0/24 is subnetted, 1 subnets
O       100.1.1.0 [110/129] via 10.1.23.2, 00:03:11, Serial0/0.32
     111.0.0.0/24 is subnetted, 1 subnets
O       111.1.1.0 [110/129] via 10.1.23.2, 00:03:11, Serial0/0.32
     10.0.0.0/24 is subnetted, 5 subnets
O       10.1.12.0 [110/128] via 10.1.23.2, 00:03:11, Serial0/0.32
O IA    10.1.45.0 [110/128] via 10.1.34.4, 00:03:11, Serial0/0.34
     11.0.0.0/24 is subnetted, 1 subnets
O       11.1.1.0 [110/129] via 10.1.23.2, 00:03:11, Serial0/0.32
```

Note R3 sees network 1.1.1.0 /24 through R2; this is because the bandwidth of

the S0/1 interface connecting R3 to R1 is 128Kbps.
Therefore, prefix 1.1.1.0 /24 is ONLY filtered from the routing table of R2 and NOT the database. The output of the following command reveals that this prefix is still in the database of R2:

On R2

R2#Show ip ospf database router 1.1.1.1 | inc Network/subnet

```
(Link ID) Network/subnet number: 111.1.1.0
(Link ID) Network/subnet number: 100.1.1.0
(Link ID) Network/subnet number: 11.1.1.0
(Link ID) Network/subnet number: 1.1.1.0
(Link ID) Network/subnet number: 10.1.13.0
(Link ID) Network/subnet number: 10.1.12.0
```

Note configuring a "distribute-list out" on R1 will NOT work at all, no other OSPF filtering solution will work except the one used in this task.

Task 7

Configure filtering on the appropriate router/s such that the existing and future routers in area 2 do NOT receive network 11.1.1.0 /24 in their routing table or their database.

The following method ONLY works for filtering LSA type 3s, and LSA type 3s ONLY. The first step is to configure a prefix-list to deny the route:

On R4

```
R4(config)#ip prefix-list TST seq 5 deny 11.1.1.0/24
R4(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32
```

Once the prefix-list is configured, it can be applied to the area that it must be filtered from, in this case area 2. This command must be configured on an ABR. In the following configuration, the prefix-list filter's network 11.1.1.0 /24 from getting IN area 2.

```
R4(config)#router ospf 1
R4(config-router)#area 2 filter-list prefix TST in
```

To verify the configuration:

On R5

R5#sh ip route ospf | inc O IA

```
O IA    1.1.1.0 [110/257] via 10.1.45.4, 00:32:51, Serial0/0.54
O IA    100.1.1.0 [110/257] via 10.1.45.4, 00:32:51, Serial0/0.54
O IA    111.1.1.0 [110/257] via 10.1.45.4, 00:32:51, Serial0/0.54
O IA    10.1.13.0 [110/909] via 10.1.45.4, 00:44:42, Serial0/0.54
O IA    10.1.12.0 [110/256] via 10.1.45.4, 00:33:01, Serial0/0.54
O IA    10.1.23.0 [110/192] via 10.1.45.4, 00:44:42, Serial0/0.54
O IA    10.1.34.0 [110/128] via 10.1.45.4, 00:44:42, Serial0/0.54
```

Note the above "Show" command reveals that R5 does NOT have the route in it's routing table, and the following command verifies that R5 does Not have the prefix in it's database.

R5#Show ip ospf database **summary** 11.1.1.0

```
          OSPF Router with ID (5.5.5.5) (Process ID 1)
```

On R4

R4#Show ip route ospf | inc O IA

```
O IA    1.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    100.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    111.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    10.1.13.0 [110/845] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    10.1.12.0 [110/192] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    10.1.23.0 [110/128] via 10.1.34.3, 00:13:32, Serial0/0.43
O IA    11.1.1.0 [110/193] via 10.1.34.3, 00:13:32, Serial0/0.43
```

Note even though the output of the above "Show" command reveals that network 11.1.1.0 /24 is in R4's routing table, the output of the following "Show" command clearly shows that it's in the database of area 0 and NOT in the database that belongs to area 2.

R4#Show ip ospf database summary 11.1.1.0

```
          OSPF Router with ID (4.4.4.4) (Process ID 1)

               Summary Net Link States (Area 0)
```

```
Routing Bit Set on this LSA
LS age: 267
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 11.1.1.0 (summary Network Number)
Advertising Router: 3.3.3.3
LS Seq Number: 80000001
Checksum: 0x950C
Length: 28
Network Mask: /24
    TOS: 0  Metric: 129
```

Task 8

Configure the appropriate router/s such that the routers in area 0 do not see network 11.1.1.0 /24 in their routing table or Link state database. You should use the same solution as the one in the previous task, but it should be implemented in the OUT bound direction.

The following method is used for filtering LSA type 3s, and LSA type 3s ONLY. Once again a prefix-list is configured to deny network 11.1.1.0 /24 on the ABR, but in the following case the "area filter-list" command is filtering network 11.1.1.0 /24 as it's advertised OUT of area 1.

On R3

```
R3(config)#ip prefix-list TST seq 5 deny 11.1.1.0/24
R3(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R3(config)#router ospf 1
R3(config-router)#area 1 filter-list prefix TST out
```

To verify the configuration:

Note the output of the following commands show that prefix 11.1.1.0 /24 is no longer in the routing table of R4 or R5.

On R4

```
R4#Show ip route 11.1.1.0 255.255.255.0
```

% Network not in table

On R5

R5#Show ip route 11.1.1.0 255.255.255.0

% Network not in table

Note the prefix is still in the routing table of R3, where the filtering is performed, but the prefix is in the routing table of this router as a route from area 1 (LSA type 1) and NOT a prefix from area 0 (From area 0's perspective this prefix is LSA type 3).

On R3

R3#Show ip route ospf

```
     1.0.0.0/24 is subnetted, 1 subnets
O       1.1.1.0 [110/129] via 10.1.23.2, 00:04:00, Serial0/0.32
     100.0.0.0/24 is subnetted, 1 subnets
O       100.1.1.0 [110/129] via 10.1.23.2, 00:04:00, Serial0/0.32
     111.0.0.0/24 is subnetted, 1 subnets
O       111.1.1.0 [110/129] via 10.1.23.2, 00:04:00, Serial0/0.32
     10.0.0.0/24 is subnetted, 5 subnets
O       10.1.12.0 [110/128] via 10.1.23.2, 00:04:00, Serial0/0.32
O IA    10.1.45.0 [110/128] via 10.1.34.4, 00:04:00, Serial0/0.34
     11.0.0.0/24 is subnetted, 1 subnets
O       11.1.1.0 [110/129] via 10.1.23.2, 00:04:00, Serial0/0.32
```

Note this is an intra-area route.

To prove this further:

On R3

R3#Sh ip ospf database summary 11.1.1.0

```
          OSPF Router with ID (3.3.3.3) (Process ID 1)
```

Note the output of the above "Show" command reveals that network 11.1.1.0 /24 is NOT in area 0, because if it was in area 0, it would have been in the Link State database of this router as a summary LSA or LSA type 3, whereas, the following "Show" command reveals that the prefix is in area 1 as a router LSA or LSA type 1.

R3#Show ip ospf database router | inc Area 1 | 11.1.1.0

```
    (Link ID) Network/subnet number: 11.1.1.0
```

Task 9

Configure the appropriate router/s such that the routers in area 0 or area 2 do not see network 111.1.1.0 /24. Use the minimum number of commands to accomplish this task.

Note the output of the following "Show" command verifies that network 111.1.1.0 /24 is in the database of R3 that belongs to Area 1, and it shows up as a router LSA or LSA type 1:

On R3

```
R3#Show ip ospf database router | inc Area 1 | 111.1.1.0

    (Link ID) Network/subnet number: 111.1.1.0
```

The output of the following "Show" command reveals that prefix 111.1.1.0 /24 is in the database of R3 as a summary LSA or LSA type 3:

```
R3#Show ip ospf database summary 111.1.1.0

            OSPF Router with ID (3.3.3.3) (Process ID 1)

                Summary Net Link States (Area 0)

  LS age: 294
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 111.1.1.0 (summary Network Number)
  Advertising Router: 3.3.3.3
  LS Seq Number: 8000000A
  Checksum: 0x6CC8
  Length: 28
  Network Mask: /24
        TOS: 0  Metric: 129
```

The following command reveals that prefix 111.1.1.0 /24 is in the routing table of R3 as an intra-area route.
Remember that intra-area routes take precedence over inter-area routes; Therefore, this prefix shows up as an "O" route in the routing table.

R3#Show ip route ospf | inc 111.1.1.0

O 111.1.1.0 [110/129] via 10.1.23.2, 00:47:21, Serial0/0.32

The following OSPF filtering mechanism works ONLY on LSA type 1s. It filters LSA type 1s from being injected into a given area, this command ONLY works if it's configured on an ABR and it is used for filtering LSA type 1 and LSA type 1s ONLY.

On R3

R3(config)#router ospf 1
R3(config-router)#**area 1 range 111.1.1.0 255.255.255.0 not-advertise**

Note the prefix is still in the routing table of R3 where the filtering is performed, but once again it shows up in the routing table as LSA type 1:

R3#Show ip route ospf | inc 111.1.1.0

O 111.1.1.0 [110/129] via 10.1.23.2, 00:00:34, Serial0/0.32

Note the prefix is no longer in the database of R3 as LSA type 3, which means that the routers in area 0 or any other area down stream to area 0 will not have this prefix in their routing table or link state database.

R3#Show ip ospf database **summary** 111.1.1.0

OSPF Router with ID (3.3.3.3) (Process ID 1)

R3#Show ip ospf database **router** | inc **Area 1** | 111.1.1.0

(Link ID) Network/subnet number: 111.1.1.0

Task 10

Configure the appropriate router/s such that none of the routers except R1 see network 100.1.1.0 /24 in their routing table; DO NOT stop advertising this network to accomplish this task. You should NOT use the solution that was used in tasks 7, 8 or 9 to accomplish this task.

On R2

You should always display the existing access-list/s and distribute-list/s before configuring one. You do not want to override an existing access-list/distribute-list few minutes before the end of your lab exam.

R2#Show access-list

Standard IP access list 1
    10 deny   1.1.1.0, wildcard bits 0.0.0.255 (19 matches)
    20 permit any (144 matches)

R2#Sh run | S router ospf 1

router ospf 1
 router-id 2.2.2.2
 log-adjacency-changes
 network 10.1.12.2 0.0.0.0 area 1
 network 10.1.23.2 0.0.0.0 area 1
 distribute-list 1 in

Note the above "Show" command verifies that there is already a distribute-list configured in the sub-router configuration mode, therefore, we should try to modify the existing access-list that is applied by the existing distribute-list.

R2(config)#**NO** access-list 1

R2(config)#access-list 1 deny 1.1.1.0 0.0.0.255
R2(config)#access-list 1 deny 100.1.1.0 0.0.0.255
R2(config)#access-list 1 permit any

To verify the configuration:

On R2

R2(config-router)#do show ip route 100.1.1.0

% Network not in table

On R3

R3#Show access-list

R3#

R3#Show ip route | inc 100.1.1.0

O   100.1.1.0 [110/129] via 10.1.23.2, 00:16:48, Serial0/0.32

R3(config)#access-list 1 deny 100.1.1.0 0.0.0.255
R3(config)#access-list 1 permit any

R3(config)#Router ospf 1
R3(config-router)#**distribute-list 1 in**

R3#Show ip route ospf | inc 100.1.1.0

R3#

On R4

R4#Show access-list

R4#

R4#Show ip route ospf | inc 100.1.1.0

O IA   100.1.1.0 [110/193] via 10.1.34.3, 05:10:53, Serial0/0.43

R4(config)#access-list 1 deny 100.1.1.0 0.0.0.255
R4(config)#access-list 1 permit any

R4(config)#router ospf 1
R4(config-router)#**distribute-list 1 in**

R4#Show ip route ospf | inc 100.1.1.0

R4#

On R5

R5#show ip route 100.1.1.0 255.255.255.0

% Network not in table

R5#Show ip ospf da summ 100.1.1.0

OSPF Router with ID (5.5.5.5) (Process ID 1)

****

Note using the "distribute-list in" sub-router configuration command ONLY effects the router that it's configured on, and the ONLY exception is if the prefix that is being filtered, is coming from area 0, meaning it's being filtered from area 0 into another area in which case it will filter the route from the database and as a result of that the routers in the non-zero area will NOT have the route in their database or routing table.
Whereas, If it's being filtered from a non-zero area into area 0, it will ONLY effect the router that it's configured on.

****

Task 11

Configure the following Loopback interfaces on R5 and redistribute these Loopback interfaces in OSPF routing domain using the default cost.

Loopback 0 – 5.5.5.5 /24
Loopback 1 – 50.5.5.5 /24
Loopback 2 – 55.5.5.5 /24

On R5

R5(config)#int lo0
R5(config-if)#ip addr 5.5.5.5 255.255.255.0

R5(config-if)#int lo1
R5(config-if)#ip addr 50.5.5.5 255.255.255.0

R5(config-if)#int lo2
R5(config-if)#ip addr 55.5.5.5 255.255.255.0

R5(config)#route-map **TST** permit 10
R5(config-route-map)#match interface lo0 lo1 lo2

R5(config)#router ospf 1
R5(config-router)#**redistribute connected subnets route-map TST**

To verify the configuration:

On R4

R4#Show ip route ospf | Inc E2

O E2   50.5.5.0 [110/20] via 10.1.45.5, 00:01:05, Serial0/0.45
O E2   55.5.5.0 [110/20] via 10.1.45.5, 00:01:05, Serial0/0.45
O E2   5.5.5.0   [110/20] via 10.1.45.5, 00:01:05, Serial0/0.45

Task 12

Configure the appropriate router such that none of the routers except R5 can see network 5.5.5.0 /24 in their routing table.

On R5

R5(config)#access-list 1 deny 5.5.5.0
R5(config)#access-list 1 permit any

R5(config)#router ospf 1
R5(config-router)#**distribute-list** 1 out

To verify the configuration:

On R4

R4#Show ip route ospf | Inc E2

O E2   50.5.5.0 [110/20] via 10.1.45.5, 00:06:00, Serial0/0.45
O E2   55.5.5.0 [110/20] via 10.1.45.5, 00:06:00, Serial0/0.45

On R1

R1#sh ip route ospf | Inc E2

O E2   50.5.5.0 [110/20] via 10.1.12.2, 00:07:08, Serial0/0.12
O E2   55.5.5.0 [110/20] via 10.1.12.2, 00:07:08, Serial0/0.12

Note this is the ONLY scenario where the "distribute-list OUT" command works in OSPF. This command MUST be configured on the ASBR or else it will not have any effect whatsoever. This command filters LSA type 5s or 7s, in this case the specific LSA type 5 is filtered from R5's OSPF database and as a result of that, none of the other OSPF routers will see the route in their routing table or database.

```
R5#Show ip ospf database external

            OSPF Router with ID (5.5.5.5) (Process ID 1)

                Type-5 AS External Link States

  LS age: 664
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 50.5.5.0 (External Network Number)
  Advertising Router: 5.5.5.5
  LS Seq Number: 80000001
  Checksum: 0x51FE
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 0

  LS age: 664
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 55.5.5.0 (External Network Number )
  Advertising Router: 5.5.5.5
  LS Seq Number: 80000001

  Checksum: 0x103B
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 0
```

Task 13

Configure the appropriate router such that none of the routers except R5 can see network 50.5.5.0 /24 in their routing table or database. You should NOT use the solution that was implemented in the previous task.

The following command is used to filter LSA type 5s or 7s, this command must be configured on an ASBR, and when configured, it filters the specified prefix from the OSPF Link state database of the ASBR.

On R5

R5(config)#router ospf 1
R5(config-router)#summary-address 50.5.5.0 255.255.255.0 **not-advertise**

Note network 50.5.5.0 /24 is NOT in the link state database of R5.

R5#Sh ip ospf da external | Inc 50.5.5.0

R5#

To verify the configuration:

On R1

R1#Sh ip route ospf | Inc E2

O E2   55.5.5.0 [110/20] via 10.1.12.2, 00:21:25, Serial0/0.12

On R4

R4#Show ip route ospf | Inc E2

O E2   55.5.5.0 [110/20] via 10.1.45.5, 00:23:17, Serial0/0.45

Task 14

Configure the appropriate router such that router R1 does NOT have network 55.5.5.0 /24 in its routing table.

Note R1 has the network in it's routing table.

On R1

R1#Show ip route ospf | inc E2

O E2   55.5.5.0 [110/20] via 10.1.12.2, 00:30:59, Serial0/0.12

Note there are no access-lists configured on this router:

R1#Sh access-list
R1#

R1(config)#access-list 1 deny 55.5.5.0
R1(config)#access-list 1 permit any

R1(config)#router ospf 1
R1(config-router)#**distribute-list 1 in**

To verify the configuration:

On R1

R1#Show ip route ospf | inc E2

Note the above "Show" command verifies that prefix 55.5.5.0 /24 was filtered successfully.

Task 15

Remove all the filters applied in the previous tasks (6 – 10, 12 - 14), if this configuration is performed successfully, all the routers should have every route advertised and redistributed in this lab.

On R1

R1(config)#**NO access-list 1**

R1(config)#router ospf 1
R1(config-router)#**NO distribute-list 1 in**

On R2

R2(config)#NO access-list 1

R2(config)#router ospf 1
R2(config-router)#NO distribute-list 1 in

On R3

R3(config)#NO access-list 1

R3(config)#router ospf 1
R3(config-router)#NO distribute-list 1 in
R3(config-router)#NO area 1 range 111.1.1.0 255.255.255.0 not-advertise
R3(config-router)#NO area 1 filter-list prefix TST out

R3(config)#NO ip prefix-list TST

On R4

R4(config)#NO access-list 1

R4(config)#router ospf 1
R4(config-router)#NO area 2 filter-list prefix TST in
R4(config-router)#NO distribute-list 1 in

R4(config)#NO ip prefix-list TST

On R5

R5(config)#NO access-list 1

R5(config)#router ospf 1
R5(config-router)#NO summary-address 50.5.5.0 255.255.255.0 not-advertise
R5(config-router)#NO distribute-list 1 out

Task 16

Configure the following loopback interfaces and advertise them in OSPF routing domain based on the following chart: These loopback interfaces should be advertised with their correct mask.

Routers	Interface / IP address	Area
R2	Loopback 0 – 2.2.2.2 /24	1
R3	Loopback 0 – 3.3.3.3 /24 Loopback 1 – 30.3.3.3 /24	1 0
R4	Loopback 0 – 4.4.4.4 /24 Loopback 1 – 40.4.4.4 /24	2 0

On R2

```
R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.255.255.0
R2(config-if)#ip ospf net point-to-point

R2(config-if)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 1
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc 2.2.2.0

    2.2.2.0 [110/65] via 10.1.12.2, 00:01:50, Serial0/0.12
```

On R3

```
R3(config)#int lo0
R3(config-if)#ip addr 3.3.3.3 255.255.255.0
R3(config-if)#ip ospf net point-to-point

R3(config-if)#int lo1
R3(config-if)#ip addr 30.3.3.3 255.255.255.0
R3(config-if)#ip ospf net point-to-point

R3(config-if)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 1
R3(config-router)#netw 30.3.3.3 0.0.0.0 area 0
```

To verify the configuration:

On R1

R1#Show ip route ospf | Inc O

```
O E2    50.5.5.0 [110/20] via 10.1.12.2, 00:01:32, Serial0/0.12
O       2.2.2.0 [110/65] via 10.1.12.2, 00:01:42, Serial0/0.12
O       3.3.3.0 [110/129] via 10.1.12.2, 00:01:42, Serial0/0.12
O E2    55.5.5.0 [110/20] via 10.1.12.2, 00:01:32, Serial0/0.12
O E2    5.5.5.0 [110/20] via 10.1.12.2, 00:01:32, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:01:42, Serial0/0.12
O IA    30.3.3.0 [110/129] via 10.1.12.2, 00:01:32, Serial0/0.12
```

On R4

```
R4(config)#int lo0
R4(config-if)#ip addr 4.4.4.4 255.255.255.0
R4(config-if)#ip ospf net point-to-point

R4(config-if)#int lo1
R4(config-if)#ip addr 40.4.4.4 255.255.255.0
R4(config-if)#ip ospf net point-to-point

R4(config-if)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 2
R4(config-router)#netw 40.4.4.4 0.0.0.0 area 0
```

To verify the configuration:

On R1

R1#Show ip route ospf | Inc O

```
O E2    50.5.5.0 [110/20] via 10.1.12.2, 00:01:47, Serial0/0.12
O       2.2.2.0 [110/65] via 10.1.12.2, 00:06:42, Serial0/0.12
O       3.3.3.0 [110/129] via 10.1.12.2, 00:06:42, Serial0/0.12
O IA    4.4.4.0 [110/193] via 10.1.12.2, 00:02:03, Serial0/0.12
O E2    55.5.5.0 [110/20] via 10.1.12.2, 00:01:48, Serial0/0.12
O E2    5.5.5.0 [110/20] via 10.1.12.2, 00:01:48, Serial0/0.12
O IA    40.4.4.0 [110/193] via 10.1.12.2, 00:01:57, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:06:42, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:06:42, Serial0/0.12
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:06:42, Serial0/0.12
O IA    30.3.3.0 [110/129] via 10.1.12.2, 00:06:32, Serial0/0.12
```

Task 17

Configure the appropriate router/s such that the routers in area 2 do NOT see any of the networks advertised by any of the routers in this topology, but routers R1, R2 and R3 and R4 should see all the networks advertised by the existing and future router/s in area 2.

By default all outgoing LSAs are flooded to the interface. This command prevents flooding of ALL OSPF LSAs out of a given interface, in this case S0/0.45.

On R4

R4(config)#int S0/0.45
R4(config-subif)#ip ospf database-filter all out

For this filtering mechanism to work, the OSPF process must be cleared.

On R5

R5#clc ip ospf proc
Reset ALL OSPF processes? [no]: y

Note R4 and R5 are still maintaining their neighbor adjacency

R5#sh ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:30	10.1.45.4	Serial0/0.54

Note R5 does NOT have any of the routes from the other routers, this includes R4 which is in the same area.

R5#Show ip route | b Gateway

Gateway of last resort is not set

```
     50.0.0.0/24 is subnetted, 1 subnets
C       50.5.5.0 is directly connected, Loopback1
     55.0.0.0/24 is subnetted, 1 subnets
C       55.5.5.0 is directly connected, Loopback2
     5.0.0.0/24 is subnetted, 1 subnets
C       5.5.5.0 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/0.54
```

Note R1 has all the routes including the ones advertised by R5.

On R1

R1#Sh ip route ospf

```
     50.0.0.0/24 is subnetted, 1 subnets
O E2    50.5.5.0 [110/20] via 10.1.12.2, 00:15:32, Serial0/0.12
     2.0.0.0/24 is subnetted, 1 subnets
O       2.2.2.0 [110/65] via 10.1.12.2, 00:28:46, Serial0/0.12
     3.0.0.0/24 is subnetted, 1 subnets
O       3.3.3.0 [110/129] via 10.1.12.2, 00:28:46, Serial0/0.12
     4.0.0.0/24 is subnetted, 1 subnets
O IA    4.4.4.0 [110/193] via 10.1.12.2, 00:28:45, Serial0/0.12
     55.0.0.0/24 is subnetted, 1 subnets
O E2    55.5.5.0 [110/20] via 10.1.12.2, 00:15:32, Serial0/0.12
     5.0.0.0/24 is subnetted, 1 subnets
O E2    5.5.5.0 [110/20] via 10.1.12.2, 00:15:32, Serial0/0.12
     40.0.0.0/24 is subnetted, 1 subnets
O IA    40.4.4.0 [110/193] via 10.1.12.2, 00:28:45, Serial0/0.12
     10.0.0.0/24 is subnetted, 5 subnets
O       10.1.23.0 [110/128] via 10.1.12.2, 00:28:46, Serial0/0.12
O IA    10.1.45.0 [110/256] via 10.1.12.2, 00:28:37, Serial0/0.12
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:28:46, Serial0/0.12
     30.0.0.0/24 is subnetted, 1 subnets
O IA    30.3.3.0 [110/129] via 10.1.12.2, 00:28:46, Serial0/0.12
```

Task 18

Configure the appropriate router/s such that the routers R1, R2 and R3 see all the routes advertised and/or redistributed by the routers in this routing domain, whereas, routers R4 ONLY see the routes advertised within their area, area2.

The "Neighbor database-filter all out" sub-router configuration command prevents flooding of ALL OSPF LSAs to a given neighbor that is reachable through an interface that has a point-to-multipoint network type at a given IP address, in this case the neighbor with an IP address of 10.1.34.4.

On R3

R3(config)#router ospf 1

R3(config-router)#neighbor 10.1.34.4 database-filter all out

Note you should get the following error message, because the above "Neighbor" command ONLY works for a neighbor through an interface that has a Point-to-Multipoint and/or NBMA OSPF network types.

%OSPF-4-CFG_NBR_INVAL_NET_TYPE: Can not use configured neighbor: neighbor command is allowed only on NBMA and point-to-multipoint networks

R3(config-router)#int S0/0.34
R3(config-subif)#ip ospf network point-to-multipoint

The same network type should be configured on R4's S0/0.43 interface, as follows:

R4(config)#int S0/0.43
R4(config-subif)#ip ospf network point-to-multipoint

R3(config-subif)#router ospf 1
R3(config-router)#neighbor 10.1.34.4 database-filter all out

Once again the OSPF process needs to be cleared:

On R4

R4#clear ip ospf proc

Reset ALL OSPF processes? [no]: y

To verify the configuration:

On R4

R4#Show ip route ospf

```
     50.0.0.0/24 is subnetted, 1 subnets
O E2    50.5.5.0 [110/20] via 10.1.45.5, 00:12:46, Serial0/0.45
     55.0.0.0/24 is subnetted, 1 subnets
O E2    55.5.5.0 [110/20] via 10.1.45.5, 00:12:46, Serial0/0.45
     5.0.0.0/24 is subnetted, 1 subnets
O E2    5.5.5.0 [110/20] via 10.1.45.5, 00:12:46, Serial0/0.45
```

On R5

R5#Sh ip route | b Gateway

```
Gateway of last resort is not set

     50.0.0.0/24 is subnetted, 1 subnets
C       50.5.5.0 is directly connected, Loopback1
     55.0.0.0/24 is subnetted, 1 subnets
C       55.5.5.0 is directly connected, Loopback2
     5.0.0.0/24 is subnetted, 1 subnets
C       5.5.5.0 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/0.54
```

Note R5 will NOT have any of the networks in its routing table because of the configuration performed in task 16.

Task 19

Erase the startup configuration of the routers and reload them before proceeding to the next lab.

Lab 9
Additional OSPF Filtering

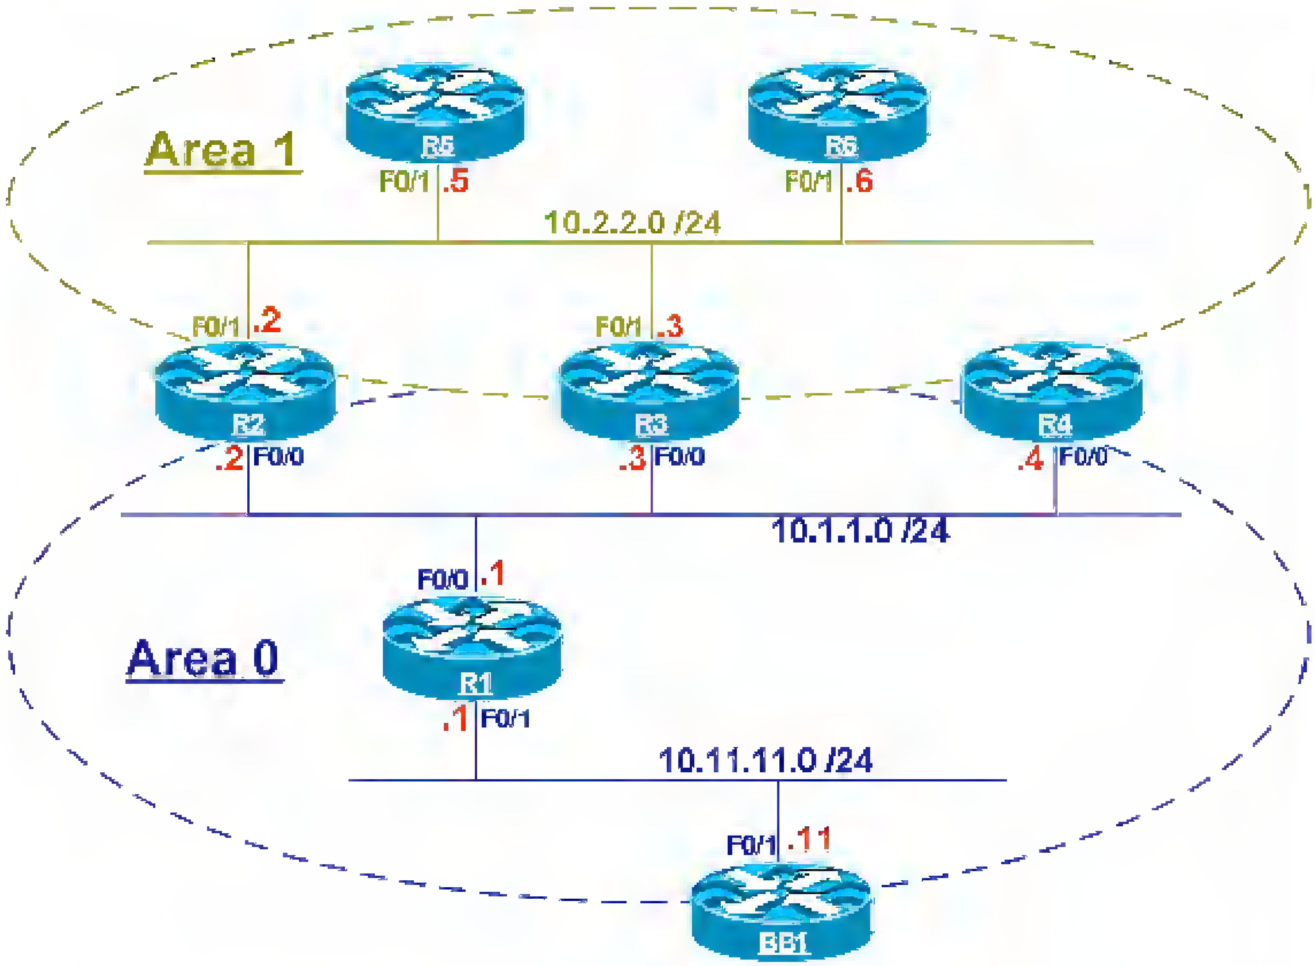

Lab Setup:

- Configure the F0/1 interface of R2, R3, R5 and R6 should be configured in VLAN 100.
- Configure the F0/0 interface of R1, R2, R3 and R4 in VLAN 200.
- Configure the F0/1 interface of R1 and BB1 in VLAN 300.
- Configure the IP addressing based on the above diagram.

Task 1

Configure the F0/1 interface of R1 and BB1 in OSPF area 0.

On both routers

(config)#Router ospf 1
(config-router)#netw 0.0.0.0 0.0.0.0 area 0

To verify the configuration:

On BB1

BB1#Ping 10.11.11.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.11.11.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

BB1#Sh ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.1.1	1	FULL/DR	00:00:38	10.11.11.1	FastEthernet0/1

Task 2

Configure R1, R2, R3 and R4's F0/0 interface in OSPF area 0.

On R2

R2(config)#router ospf 1
R2(config-router)#network 10.1.1.2 0.0.0.0 area 0

On R3

R3(config)#router ospf 1
R3(config-router)#network 10.1.1.3 0.0.0.0 area 0

On R4

```
R4(config)#router ospf 1
R4(config-router)#network 10.1.1.4 0.0.0.0 area 0
```

To verify the configuration:

On R1

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.11.11.11       1   FULL/BDR        00:00:34    10.11.11.11     FastEthernet0/1
10.1.1.4          1   FULL/DROTHER    00:00:36    10.1.1.4        FastEthernet0/0
10.2.2.2          1   FULL/BDR        00:00:30    10.1.1.2        FastEthernet0/0
10.2.2.3          1   FULL/DROTHER    00:00:31    10.1.1.3        FastEthernet0/0
```

Task 3

Configure the F0/1 interface of R2, R3, R5 and R6 in Area 1.

On R2

```
R2(config)#router ospf 1
R2(config-router)#network 10.2.2.2 0.0.0.0 area 1
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#network 10.2.2.3 0.0.0.0 area 1
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#network 10.2.2.5 0.0.0.0 area 1
```

On R6

```
R6(config)#router ospf 1
```

R6(config-router)#network 10.2.2.6 0.0.0.0 area 1

To verify the configuration:

On R5

R5#Show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.2.2.6	1	FULL/BDR	00:00:39	10.2.2.6	FastEthernet0/1
10.2.2.2	1	FULL/DROTHER	00:00:33	10.2.2.2	FastEthernet0/1
10.2.2.3	1	FULL/DROTHER	00:00:33	10.2.2.3	FastEthernet0/1

On R6

R6#Show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.2.2.5	1	FULL/DR	00:00:38	10.2.2.5	FastEthernet0/1
10.2.2.2	1	FULL/DROTHER	00:00:34	10.2.2.2	FastEthernet0/1
10.2.2.3	1	FULL/DROTHER	00:00:35	10.2.2.3	FastEthernet0/1

Task 4

Configure two loopback interfaces on R5 using the following IP addresses:

Lo0 = 5.5.5.5 /8 and Lo1 = 55.5.5.5 /24
Lo1 interface should be advertised in OSPF area 1 with its correct mask.
Lo0 should be redistributed in OSPF as metric-type 1.

On R5

R5(config)#inter lo0
R5(config-if)#ip address 5.5.5.5 255.0.0.0

R5(config)#inter lo1
R5(config-if)#ip address 55.5.5.5 255.255.255.0
R5(config-if)#ip ospf network point-to-point

R5(config)#route-map **TST** permit 10

```
R5(config)#match interface lo0

R5(config)#router ospf 1
R5(config-router)#network 55.5.5.5 0.0.0.0 area 1
R5(config-router)#redistribute connected metric-type 1 subnets route-map TST
```

Task 5

Configure two loopback interfaces on R6 using the following IP addresses:

Lo0 = 6.6.6.6 /8 and Lo1 = 66.6.6.6 /24
Lo1 interface should be advertised in OSPF area 1 with its correct mask.
Lo0 should be redistributed in OSPF as metric-type 1.

On R6

```
R6(config)#inter lo0
R6(config-if)#ip address 6.6.6.6 255.0.0.0

R6(config)#inter lo1
R6(config-if)#ip address 66.6.6.6 255.255.255.0
R6(config-if)#ip ospf network point-to-point

R6(config)#route-map TST permit 10
R6(config)#match interface lo0

R6(config)#router ospf 1
R6(config-router)#network 66.6.6.6 0.0.0.0 area 1
R6(config-router)#redistribute connected metric-type 1 subnets route-map TST
```

Task 6

Configure two loopback interfaces on R4 using the following IP addresses:

Lo0 = 4.4.4.4 /8 and Lo1 = 44.4.4.4 /24
Lo0 interface should be advertised in OSPF area 1 with its correct mask.
Lo1 should be redistributed in OSPF as metric-type 2, this route should be tagged with 44 as it gets redistributed in OSPF routing domain.

On R4

```
R4(config)#interface Loopback0
R4(config-if)#ip address 4.4.4.4 255.0.0.0
R4(config-if)#ip ospf network point-to-point

R4(config)#interface Loopback1
R4(config-if)#ip address 44.4.4.4 255.255.255.0

R4(config)#route-map TST permit 10
R4(config-route-map)#match interface Loopback1
R4(config-route-map)#set tag 44

R4(config)#router ospf 1
R4(config-router)#redistribute connected subnets route-map TST
R4(config-router)#network 4.4.4.4 0.0.0.0 area 1
R4(config-router)#network 10.1.1.4 0.0.0.0 area 0
```

To verify the configuration:

On R1

R1#Sh ip route ospf

```
O IA 4.0.0.0/8     [110/2] via 10.1.1.4, 00:49:17, FastEthernet0/0
O IA 55.5.5.0/24 [110/3] via 10.1.1.3, 00:49:17, FastEthernet0/0
                        [110/3] via 10.1.1.2, 00:49:17, FastEthernet0/0
O E1 5.0.0.0/8 [110/22] via 10.1.1.3, 00:49:17, FastEthernet0/0
                        [110/22] via 10.1.1.2, 00:49:17, FastEthernet0/0
     66.0.0.0/24 is subnetted, 1 subnets
O IA    66.6.6.0 [110/3] via 10.1.1.3, 00:49:17, FastEthernet0/0
                        [110/3] via 10.1.1.2, 00:49:17, FastEthernet0/0
O E2 6.0.0.0/8 [110/20] via 10.1.1.3, 00:49:17, FastEthernet0/0
                        [110/20] via 10.1.1.2, 00:49:17, FastEthernet0/0
     10.0.0.0/24 is subnetted, 3 subnets
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:49:17, FastEthernet0/0
                        [110/2] via 10.1.1.2, 00:49:17, FastEthernet0/0
     44.0.0.0/24 is subnetted, 1 subnets
O E2    44.4.4.0 [110/20] via 10.1.1.4, 00:49:17, FastEthernet0/0
```

R1#Sh ip route 44.4.4.4

```
Routing entry for 44.4.4.0/24
  Known via "ospf 1", distance 110, metric 20
```

Tag 44, type extern 2, forward metric 1
Last update from 10.1.1.4 on FastEthernet0/0, 00:52:12 ago
Routing Descriptor Blocks:
* 10.1.1.4, from 10.1.1.4, 00:52:12 ago, via FastEthernet0/0
Route metric is 20, traffic share count is 1
Route tag 44

On BB1

BB1#Sh ip route ospf

O IA 4.0.0.0/8 [110/3] via 10.11.11.1, 00:53:43, FastEthernet0/1
O IA 55.0.0.0/8 [110/4] via 10.11.11.1, 00:53:43, FastEthernet0/1
O E1 5.0.0.0/8 [110/23] via 10.11.11.1, 00:53:43, FastEthernet0/1
66.0.0.0/24 is subnetted, 1 subnets
O IA 66.6.6.6 [110/4] via 10.11.11.1, 00:53:43, FastEthernet0/1
O E2 6.0.0.0/8 [110/20] via 10.11.11.1, 00:53:43, FastEthernet0/1
10.0.0.0/24 is subnetted, 3 subnets
O IA 10.2.2.0 [110/3] via 10.11.11.1, 00:53:43, FastEthernet0/1
O 10.1.1.0 [110/2] via 10.11.11.1, 00:53:43, FastEthernet0/1
44.0.0.0/24 is subnetted, 1 subnets
O E2 44.4.4.0 [110/20] via 10.11.11.1, 00:53:43, FastEthernet0/1

Task 7

Configure R1 to filter all networks that carry a tag of 44.

On R1

R1(config)#route-map **Task-7** deny 10
R1(config-router)#match tag 44

R1(config)#route-map **Task-7** **permit** 20

In the above configuration routes that carry a tag of 44 are matched and denied and the routes that do NOT carry a tag of 44 are all permitted.

R1(config)#router ospf 1
R1(config-router)#**distribute-list route-map Task-7 in**

To verify the configuration:

On R1

R1#Show ip route ospf

```
O IA 4.0.0.0/8   [110/2] via 10.1.1.4, 00:02:12, FastEthernet0/0
O IA 55.5.5.0/24 [110/3] via 10.1.1.3, 00:02:12, FastEthernet0/0
                 [110/3] via 10.1.1.2, 00:02:12, FastEthernet0/0
O E1 5.0.0.0/8   [110/22] via 10.1.1.3, 00:02:12, FastEthernet0/0
                 [110/22] via 10.1.1.2, 00:02:12, FastEthernet0/0
     66.0.0.0/24 is subnetted, 1 subnets
O IA    66.6.6.0 [110/3] via 10.1.1.3, 00:02:12, FastEthernet0/0
                 [110/3] via 10.1.1.2, 00:02:12, FastEthernet0/0
O E2 6.0.0.0/8   [110/20] via 10.1.1.3, 00:02:12, FastEthernet0/0
                 [110/20] via 10.1.1.2, 00:02:12, FastEthernet0/0
     10.0.0.0/24 is subnetted, 3 subnets
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:02:12, FastEthernet0/0
                 [110/2] via 10.1.1.2, 00:02:12, FastEthernet0/0
```

Note network 44.4.4.0 /24 is blocked because it carried a tag of 44. Its very important to note that this filtering is performed on R1 and the routes that carry a tag of 44 are filtered from the routing table of R1 and NOT the database.

To verify the configuration:

On BB1

BB1#Sh ip route ospf

```
O IA 4.0.0.0/8   [110/3] via 10.11.11.1, 00:08:42, FastEthernet0/1
O IA 55.5.5.0/24 [110/4] via 10.11.11.1, 00:08:42, FastEthernet0/1
O E1 5.0.0.0/8   [110/23] via 10.11.11.1, 00:08:42, FastEthernet0/1
     66.0.0.0/24 is subnetted, 1 subnets
O IA    66.6.6.0 [110/4] via 10.11.11.1, 00:08:42, FastEthernet0/1
O E2 6.0.0.0/8   [110/20] via 10.11.11.1, 00:08:42, FastEthernet0/1
     10.0.0.0/24 is subnetted, 3 subnets
O IA    10.2.2.0 [110/3] via 10.11.11.1, 00:08:42, FastEthernet0/1
O       10.1.1.0   [110/2] via 10.11.11.1, 00:08:42, FastEthernet0/1
     44.0.0.0/24 is subnetted, 1 subnets
O E2    44.4.4.0 [110/20] via 10.11.11.1, 00:08:42, FastEthernet0/1
```

On R1

R1#Sh ip ospf da ex adv-router 10.1.1.4

```
            OSPF Router with ID (10.1.1.1) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 1036
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 44.4.4.0 (External Network Number )
  Advertising Router: 10.1.1.4
  LS Seq Number: 80000004
  Checksum: 0xE04C
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 20
        Forward Address: 0.0.0.0
        External Route Tag: 44
```

Note the route is still in the database of R1.

Task 8

Remove the "distribute-list route-map Task-7 in "command from the previous task and configure R1 to filter all OSPF external type 2 prefixes. You should NOT configure an access-list to accomplish this task.

On R1

```
R1(config)#route-map Task-8 deny 10
R1(config-route-map)#match route-type external type-2

R1(config)#route-map Task-8 permit 20

R1(config)#router ospf 1
R1(config-router)#distribute-list route-map Task-8 in
```

To verify the configuration:

On R1

R1#Sh ip route ospf

```
O IA 4.0.0.0/8   [110/2] via 10.1.1.4, 00:00:09, FastEthernet0/0
O IA 55.5.5.0/24 [110/3] via 10.1.1.3, 00:00:09, FastEthernet0/0
                 [110/3] via 10.1.1.2, 00:00:09, FastEthernet0/0
O E1 5.0.0.0/8   [110/22] via 10.1.1.3, 00:00:09, FastEthernet0/0
                 [110/22] via 10.1.1.2, 00:00:09, FastEthernet0/0
     66.0.0.0/24 is subnetted, 1 subnets
O IA    66.6.6.0 [110/3] via 10.1.1.3, 00:00:09, FastEthernet0/0
                 [110/3] via 10.1.1.2, 00:00:09, FastEthernet0/0
     10.0.0.0/24 is subnetted, 3 subnets
O IA    10.2.2.0 [110/2] via 10.1.1.3, 00:00:09, FastEthernet0/0
                 [110/2] via 10.1.1.2, 00:00:09, FastEthernet0/0
```

There is no need to remove the previous distribute-list command, when a new one is entered; it overrides the previous distribute-list command.
Note the external type-2 (E2) routes are filtered from the routing table of R1, but they are still in the database of this router and therefore, as a result of that, BB1 will have the E2 routes in its routing table.

On BB1

BB1#Sh ip route ospf

```
O IA 4.0.0.0/8   [110/3] via 10.11.11.1, 00:21:51, FastEthernet0/1
O IA 55.5.5.0/24 [110/4] via 10.11.11.1, 00:21:51, FastEthernet0/1
O E1 5.0.0.0/8   [110/23] via 10.11.11.1, 00:21:51, FastEthernet0/1
     66.0.0.0/24 is subnetted, 1 subnets
O IA    66.6.6.0 [110/4] via 10.11.11.1, 00:21:51, FastEthernet0/1
O E2 6.0.0.0/8  [110/20] via 10.11.11.1, 00:21:51, FastEthernet0/1
     10.0.0.0/24 is subnetted, 3 subnets
O IA    10.2.2.0 [110/3] via 10.11.11.1, 00:21:51, FastEthernet0/1
O       10.1.1.0   [110/2] via 10.11.11.1, 00:21:51, FastEthernet0/1
     44.0.0.0/24 is subnetted, 1 subnets
O E2    44.4.4.0 [110/20] via 10.11.11.1, 00:21:51, FastEthernet0/1
```

Task 9

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 10
Redirecting traffic in OSPF

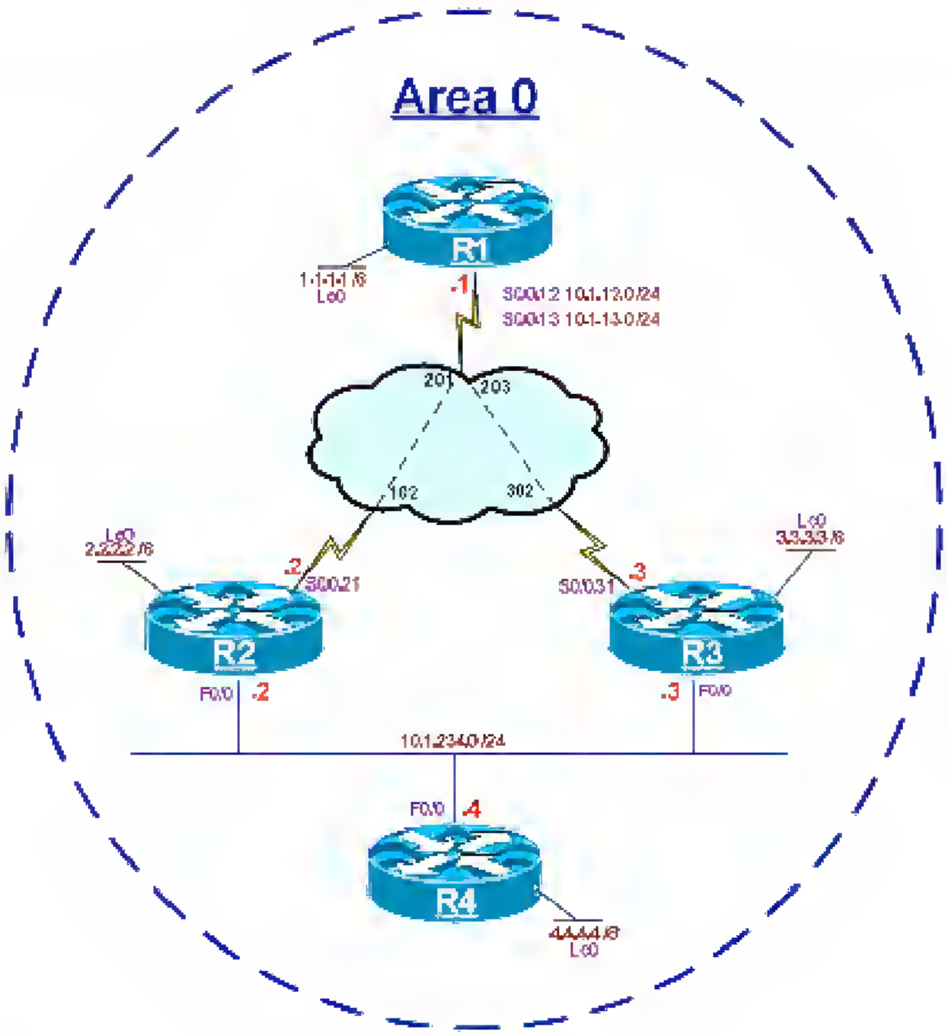

Lab Setup:

- R1 should be configured with two point-to-point links, one connecting R1 to R2 and the other connecting R1 to R3. R2 and R3 should also be configured in a point-to-point manner.
- R2, R3 and R4's F0/0 interfaces should be configured to be in VLAN 234.
- Use the IP address chart below to assign IP addresses to the routers.

IP Addressing:

Router	Interface	IP address	Area
R1	Lo0 F/R interface to R2 F/R interface to R3	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24	Area 0 Area 0 Area 0
R2	**Lo0** **F/R interface to R1** **F0/0 interface**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.234.2 /24**	**Area 0** **Area 0** **Area 0**
R3	Lo0 F/R interface to R1 F0/0 interface	3.3.3.3 /8 10.1.13.3 /24 10.1.234.3 /24	Area 0 Area 0 Area 0
R4	**Lo0** **F0/0 interface**	**4.4.4.4 /8** **10.1.234.4 /24**	**Area 0** **Area 0**

Task 1

Configure OSPF on all routers and advertise their directly connected network in area 0 and ensure that these routers can reach all the advertised networks. Ensure that the loopback interfaces are advertised with their correct mask.

On All Routers

```
(config-if)#router ospf 1
(config-router)#netw 0.0.0.0 0.0.0.0 are 0

(config-router)#interface Lo0
(config-if)#ip ospf network point-to-point
```

Task 2

R1 has two ways to reach network 4.0.0.0 /8, ensure that R1 uses R2 to reach this network. R1 should go directly to R3 to reach network 3.0.0.0.
However, if R2 goes down, R3 should be used as a transit router to reach network 4.0.0.0 /8.
DO NOT USE THE FOLLOWING COMMANDS:

Bandwidth, any global config command, OSPF cost command or the distance command.

Note the routing table of R1 reveals that R1 can reach network 4.0.0.0 via R2 and R3.

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:00:21, Serial0/0.12
O    3.0.0.0/8 [110/65] via 10.1.13.3, 00:00:21, Serial0/0.13
O    4.0.0.0/8 [110/66] via 10.1.13.3, 00:00:21, Serial0/0.13
               [110/66] via 10.1.12.2, 00:00:21, Serial0/0.12
     10.1.0.0/24 is subnetted, 3 subnets
O       10.1.234.0 [110/65] via 10.1.13.3, 00:00:21, Serial0/0.13
                   [110/65] via 10.1.12.2, 00:00:21, Serial0/0.12
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#max-metric router-lsa
```

The above command will cause a router to originate LSAs with a maximum metric of 0xffff (LSInfinity). This is done so that other routers do not prefer the router as a transit hop in their path to a given network.

To Verify the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B – BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
O    2.0.0.0/8 [110/65] via 10.1.12.2, 00:04:12, Serial0/0.12
O    3.0.0.0/8 [110/65] via 10.1.13.3, 00:04:12, Serial0/0.13
O    4.0.0.0/8 [110/66] via 10.1.12.2, 00:04:12, Serial0/0.12
     10.1.0.0/24 is subnetted, 3 subnets
O       10.1.234.0 [110/65] via 10.1.12.2, 00:04:12, Serial0/0.12
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

Task 3

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 11
Database Overload Protection

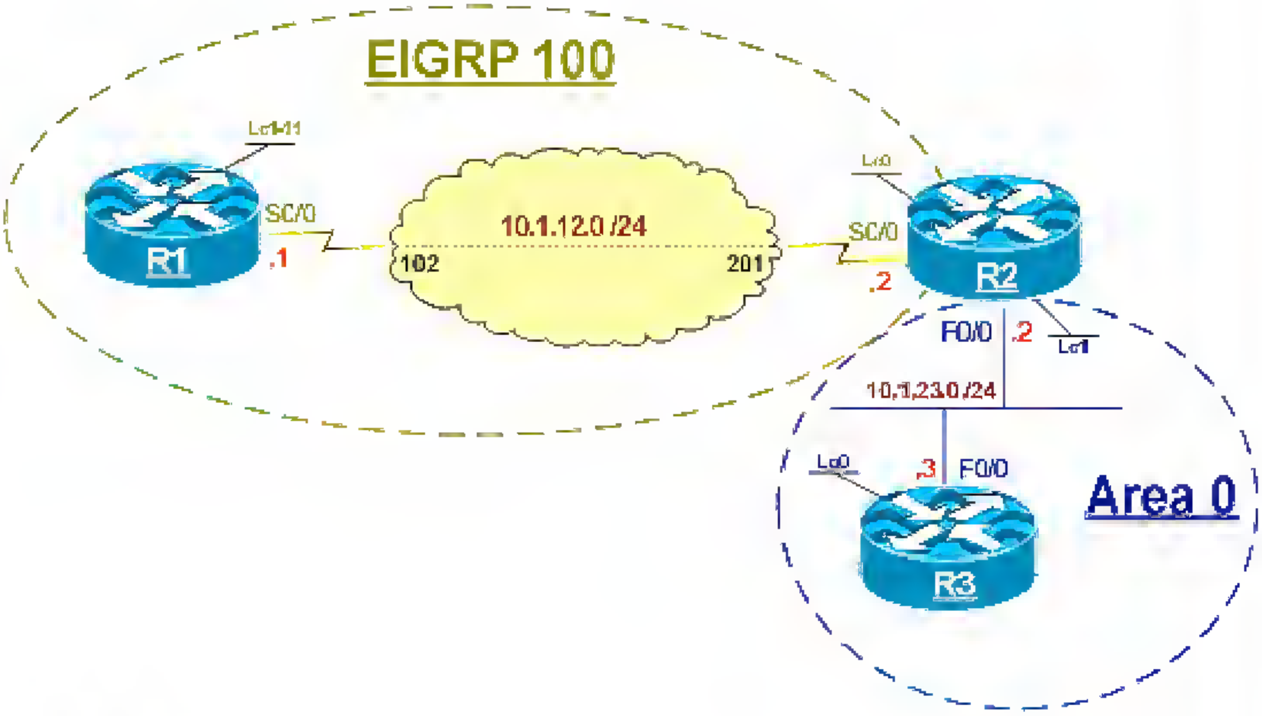

Lab Setup:

- The frame-relay connection between R1 and R2 must be configured in a point-to-point manner.

- R2 and R3's F0/0 interface should be configured in VLAN 23.

IP Addressing:

Router	Interface	IP address	Area
R1	Lo1	1.0.0.1 /8	Eigrp 100
	Lo2	2.0.0.1 /8	Eigrp 100
	Lo3	3.0.0.1 /8	Eigrp 100
	Lo4	4.0.0.1 /8	Eigrp 100
	Lo5	5.0.0.1 /8	Eigrp 100
	Lo6	6.0.0.1 /8	Eigrp 100
	Lo7	7.0.0.1 /8	Eigrp 100
	Lo8	8.0.0.1 /8	Eigrp 100
	Lo9	9.0.0.1 /8	Eigrp 100
	Lo10	10.0.0.1 /24	Eigrp 100
	Lo11	11.0.0.1 /8	Eigrp 100
	S0/0.12	10.1.12.1 /24	Eigrp 100
R2	Lo0	2.2.2.2 /8	Eigrp 100
	Lo1	22.2.2.2 /8	OSPF area 0
	S0/0.21	10.1.12.2 /24	Eigrp 100
	F0/0	10.1.23.2 /24	OSPF area 0
R3	Lo0	3.3.3.3 /8	OSPF area 0
	F0/0	10.1.23.3 /24	OSPF area 0

Task 1

Configure the routers as follows:

- On R1, configure Eigrp 100 and advertise networks 1.0.0.0 /8 – 5.0.0.0 /8 and the frame-relay interface to R2.
- On R2, configure Eigrp 100 and advertise networks 2.0.0.0 /8 and the frame-relay interface link to R1.
- On R2, configure OSPF and advertise network 22.0.0.0 /8 and the F0/0 interface to R3 in area 0.
- On R3, configure OSPF and advertise all of its interfaces in area 0.

On R1

```
R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
```

```
R1(config-router)#netw 2.0.0.0
R1(config-router)#netw 3.0.0.0
R1(config-router)#netw 4.0.0.0
R1(config-router)#netw 5.0.0.0
R1(config-router)#netw 10.1.12.0 0.0.0.255
```

On R2

```
R2(config-if)#router ospf 1
R2(config-router)#netw 22.2.2.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0

R2(config-router)#router eigrp 100
R2(config-router)#netw 10.1.12.0 0.0.0.255
R2(config-router)#netw 2.0.0.0
R2(config-router)#no auto
```

On R3

```
R3(config-if)#router ospf 1
R3(config-router)#netw 0.0.0.0 0.0.0.0 are 0
```

Task 2

On R2, perform a mutual redistribution between OSPF and Eigrp 100.

On R2

```
R2(config)#router eigrp 100
R2(config-router)#redistribute ospf 1 metric 1500 20000 255 1 1500

R2(config-router)#router ospf 1
R2(config-router)#redistribute eigrp 100 subnets
```

Note when redistributing routes into OSPF, they will be redistributed with a metric of 20.

To verify the configuration:

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

O E2 1.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
O E2 2.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
C    3.0.0.0/8 is directly connected, Loopback0
O E2 4.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
O E2 5.0.0.0/8 [110/20] via 10.1.23.2, 00:01:18, FastEthernet0/0
     22.0.0.0/32 is subnetted, 1 subnets
O       22.2.2.2 [110/2] via 10.1.23.2, 00:01:18, FastEthernet0/0
     10.1.0.0/24 is subnetted, 2 subnets
O E2    10.1.12.0 [110/20] via 10.1.23.2, 00:01:20, FastEthernet0/0
C       10.1.23.0 is directly connected, FastEthernet0/0
```

Task 3

R2 should be configured such that the maximum number of prefixes that can be redistributed into OSPF routing protocol is 10.
R2 should generate two warning messages. The first message should occur when the number of redistributed prefixes reaches 70% of the configured threshold (10). The second message should occur when the 10th prefix is redistributed.

On R2

```
R2(config-router)#router ospf 1
R2(config-router)#redistribute maximum-prefix 10 70 warning-only
```

The above command limits the number of prefixes that can be redistributed into OSPF routing domain. In this case, the router will generate two warning messages, the first one will be generated when 70% of the configured threshold (10) is reached and the second message will be generated when the configured threshold (10) is exceeded.

The initial config file has created 11 Loopback interfaces for testing purpose. Test this policy by advertising these loopback interfaces one at a time and observe the warning messages.

Task 4

The administrator of R1 is constantly violating the maximum routes policy, in order to safeguard against this, you should configure R2 such that only 10 prefixes are allowed to be redistributed into OSPF, if R1 advertises more than 10 prefixes in Eigrp 100, R2 should ignore the extra prefixes.

On R2

```
R2(config)#router ospf 1
R2(config-router)#no redistribute maximum-prefix 10 70 warning-only
R2(config-router)#redistribute maximum-prefix 10 70
```

Since the "warning-only" keyword is not used, R2 will ignore any advertisement above the set threshold.

Task 5

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 12 – OSPF Non-Broadcast Networks

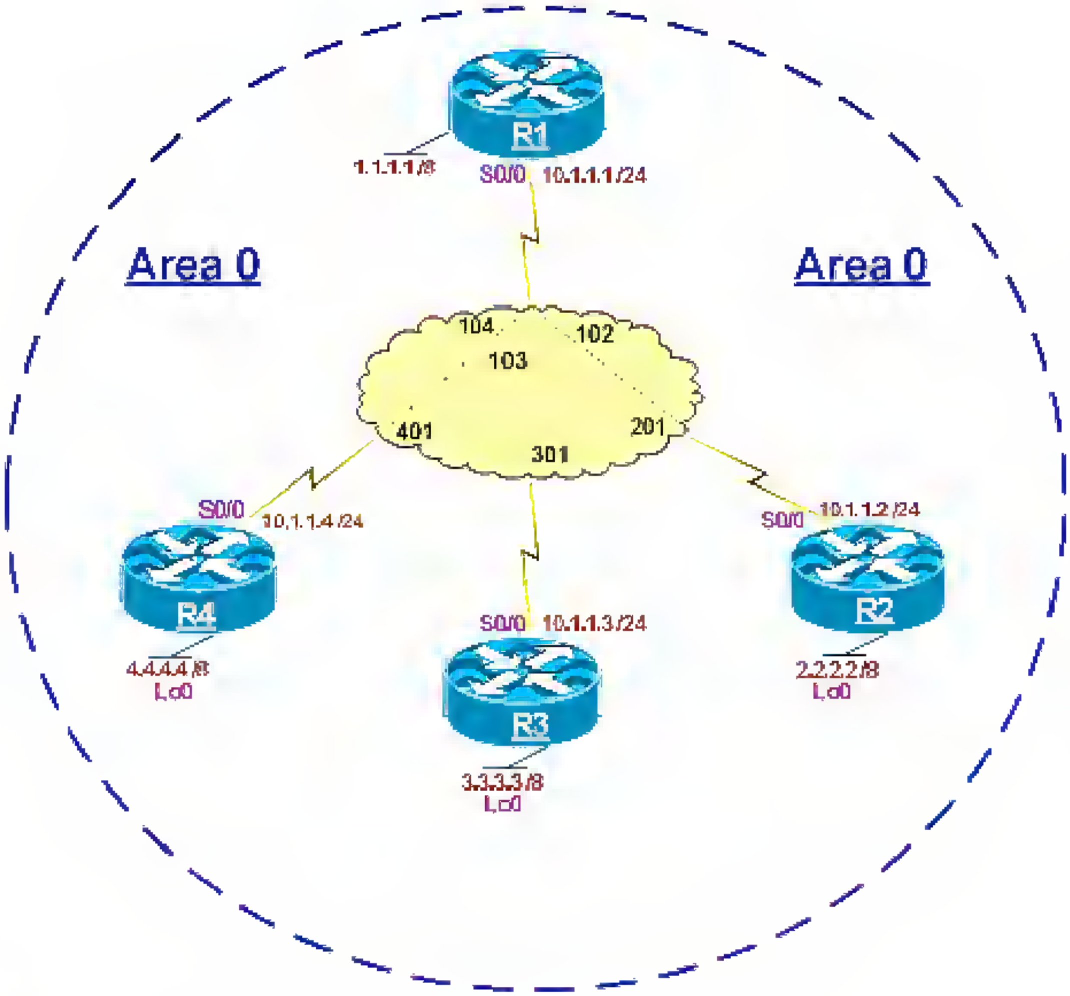

Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a Frame-relay Multipoint manner. DO NOT configure sub-interfaces on any of the routers. Use the broadcast keyword when configuring the "Frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	S0/0 = 10.1.1.2 /24 Loopback0 = 2.2.2.2 /8
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8
R4	S0/0 = 10.1.1.4 /24 Loopback0 = 4.4.4.4 /8

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in Area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. DO NOT change the network type to accomplish this task.

On R1

```
R1(config)#router ospf 1
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point
```

Note the following command is required since R1 is the hub.

```
R1(config)#interface Serial0/0
R1(config-if)#ip ospf priority 255
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0

R2(config-router)#int lo0
```

R2(config-if)#ip ospf network point-to-point

On R3

R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point

On R4

R4(config)#router ospf 1
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point

Note the following command is required so the spokes will NOT participate in DR/BDR election.

On R2, R3 and R4

(config)#interface Serial0/0
(config-if)#ip ospf priority 0

To verify the configuration:

On R1

R1#Sh ip ospf neighbor

R1#Show ip route ospf

Note there is NO neighbor adjacencies established, as a result of that; there won't be any routes in the OSPF routing table.
The reason is the OSPF network type, the default OSPF network type on Multipoint Frame-relay interface is NON_BROADCAST, the following show command reveals the OSPF network type.

On R1

R1#Show ip ospf interface S0/0

The network type

```
Serial0/0 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type NON_BROADCAST, Cost: 64
  Transmit Delay is 1 sec, State DR, Priority 255
  Designated Router (ID) 1.1.1.1, Interface address 10.1.1.1
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:22
  Supports Link-local Signaling (LLS)
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

Since the task states that the network type can not be changed, then, the "priority" sub-router configuration command can be used to accomplish this task; in this case the priority command needs to be configured on the hub router ONLY.

On R1

```
R1(config)#router ospf 1
R1(config-router)#neighbor 10.1.1.2
R1(config-router)#neighbor 10.1.1.3
R1(config-router)#neighbor 10.1.1.4
```

To verify the configuration:

On R1

R1#Show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	0	FULL/DROTHER	00:01:42	10.1.1.2	Serial0/0
3.3.3.3	0	FULL/DROTHER	00:01:48	10.1.1.3	Serial0/0
4.4.4.4	0	FULL/DROTHER	00:01:50	10.1.1.4	Serial0/0

R1#Show ip route ospf

```
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:52, Serial0/0
```

O   3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:52, Serial0/0
O   4.0.0.0/8 [110/65] via 10.1.1.4, 00:00:52, Serial0/0

Note once the "Neighbor" command is configured on the hub router, the routers will transition into FULL state and exchange routes.
Remember, when the "Neighbor" command is configured, all OSPF packets will use UNICAST instead of MULTICAST; Therefore, there the "frame-relay map" commands DO NOT need to be configured with the "Broadcast" keyword.

On R1

R1(config)#int s0/0
R1(config-if)#NO frame-relay map ip 10.1.1.2 102 broadcast
R1(config-if)#NO frame-relay map ip 10.1.1.3 103 broadcast
R1(config-if)#NO frame-relay map ip 10.1.1.4 104 broadcast

R1(config-if)#frame-relay map ip 10.1.1.2 102
R1(config-if)#frame-relay map ip 10.1.1.3 103
R1(config-if)#frame-relay map ip 10.1.1.4 104

On R2

R2(config)#int S0/0
R2(config-if)#NO frame-relay map ip 10.1.1.1 201 broadcast

R2(config-if)#frame-relay map ip 10.1.1.1 201

On R3

R3(config)#int S0/0
R3(config-if)#NO frame-relay map ip 10.1.1.1 301 broadcast

R3(config-if)#frame-relay map ip 10.1.1.1 301

On R4

R4(config)#int S0/0
R4(config-if)#NO frame-relay map ip 10.1.1.1 401 broadcast

R4(config-if)#frame-relay map ip 10.1.1.1 401

To test the configuration:

On R1

R1#Show ip route ospf

```
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:10, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:10, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:00:10, Serial0/0
```

On R2

R2#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:07:21, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:07:21, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:07:21, Serial0/0
```

On R3

R3#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:07:57, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:07:57, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:07:57, Serial0/0
```

On R4

R4#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:08:30, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:08:30, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:08:30, Serial0/0
```

Task 2

Ensure that every router has NLRI to the loopback interfaces advertised in OSPF, you should use ping to test and verify reachability.

To test the reachability:

On R2

R2#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 08:27:26, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 08:27:26, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 08:27:26, Serial0/0
```

Note R2 will not have reachability to networks 3.0.0.0/8 or 4.0.0.0/8, because it does not have layer 2 mapping for the next hop IP address.
In OSPF NON_BROADCAST network type the next hop IP address is the IP address of the router that originated the route and NOT the router that advertised it, therefore, the spokes will not have NLRI to networks advertised by other spokes, this problem can be resolved by configuring the following Frame-relay map commands:
On each spoke a "Frame-relay map" command is configured for the frame-relay interface IP address of the other spokes using their only DLCI pointing to the hub, as follows:

On R2

```
R2(config)#interface S0/0
R2(config-if)#Frame-relay map ip 10.1.1.3 201
R2(config-if)#Frame-relay map ip 10.1.1.4 201
```

On R3

```
R3(config)#interface S0/0
R3(config-if)#Frame-relay map ip 10.1.1.2 301
R3(config-if)#Frame-relay map ip 10.1.1.4 301
```

On R4

```
R4(config)#interface S0/0
R4(config-if)#Frame-relay map ip 10.1.1.2 401
R4(config-if)#Frame-relay map ip 10.1.1.3 401
```

To test the configuration:

On R2

R2#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
```

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms
```

R2#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/120 ms
```

On R3

R3#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/125/168 ms
```

R3#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

On R4

R4#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/121/148 ms
```

R4#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

Task 3

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 13 – OSPF Broadcast Networks

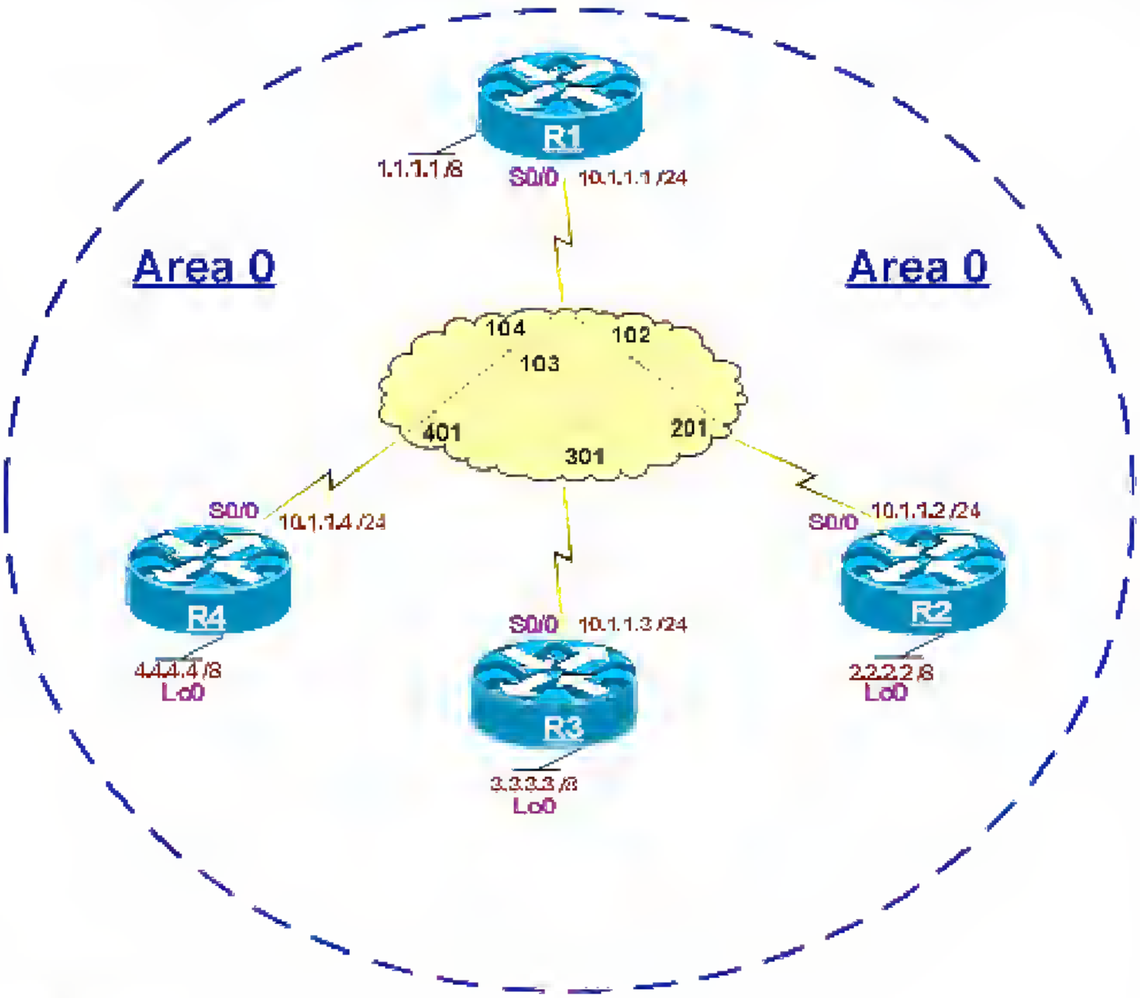

Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a Frame-relay Multipoint manner. DO NOT configure sub-interfaces on any of the routers. Use the "broadcast" keyword when configuring the "Frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	S0/0 = 10.1.1.2 /24 Loopback0 = 2.2.2.2 /8
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8
R4	S0/0 = 10.1.1.4 /24 Loopback0 = 4.4.4.4 /8

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in Area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. You should use OSPF BROADCAST network type to accomplish this task.

On R1

R1(config)#router ospf 1
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point

Note the following command is required since R1 is the hub.

R1(config)#interface Serial0/0
R1(config-if)#ip ospf network broadcast
R1(config-if)#ip ospf priority 255

On R2

R2(config)#router ospf 1
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point

On R3

R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point

On R4

R4(config)#router ospf 1
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point

Note the following command is required so the spokes will NOT participate in DR/BDR election; it also changes the network type to BROADCAST.

On R2, R3 and R4

(config)#interface Serial0/0
(config-if)#ip ospf network broadcast
(config-if)#ip ospf priority 0

To verify the configuration:

On R1

R1#Sh ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	0	FULL/DROTHER	00:00:32	10.1.1.2	Serial0/0
3.3.3.3	0	FULL/DROTHER	00:00:33	10.1.1.3	Serial0/0
4.4.4.4	0	FULL/DROTHER	00:00:31	10.1.1.4	Serial0/0

R1#Sh ip route ospf

O 2.0.0.0/8 [110/65] via 10.1.1.2, 00:01:51, Serial0/0
O 3.0.0.0/8 [110/65] via 10.1.1.3, 00:01:51, Serial0/0
O 4.0.0.0/8 [110/65] via 10.1.1.4, 00:01:51, Serial0/0

On R2

R2#Show ip route ospf

O 1.0.0.0/8 [110/65] via 10.1.1.1, 00:02:29, Serial0/0
O 3.0.0.0/8 [110/65] via 10.1.1.3, 00:02:29, Serial0/0
O 4.0.0.0/8 [110/65] via 10.1.1.4, 00:02:29, Serial0/0

Note the next hop was NOT changed
Just like the previous lab

On R3

R3#Show ip route ospf

O 1.0.0.0/8 [110/65] via 10.1.1.1, 00:03:09, Serial0/0
O 2.0.0.0/8 [110/65] via 10.1.1.2, 00:03:09, Serial0/0
O 4.0.0.0/8 [110/65] via 10.1.1.4, 00:03:09, Serial0/0

On R4

R4#Show ip route ospf

O 1.0.0.0/8 [110/65] via 10.1.1.1, 00:03:46, Serial0/0
O 2.0.0.0/8 [110/65] via 10.1.1.2, 00:03:46, Serial0/0
O 3.0.0.0/8 [110/65] via 10.1.1.3, 00:03:46, Serial0/0

Note Once again the next hop IP address is pointing to the router that advertised the route, in this case the frame-relay solution from the previous lab can also be used as the solution to this problem, but remember that the "broadcast" keyword should NOT be used when configuring the "Frame-relay map" statements on the spokes pointing to the frame-relay interface IP address of the other spokes.

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 14 – OSPF Point-to-Point Networks

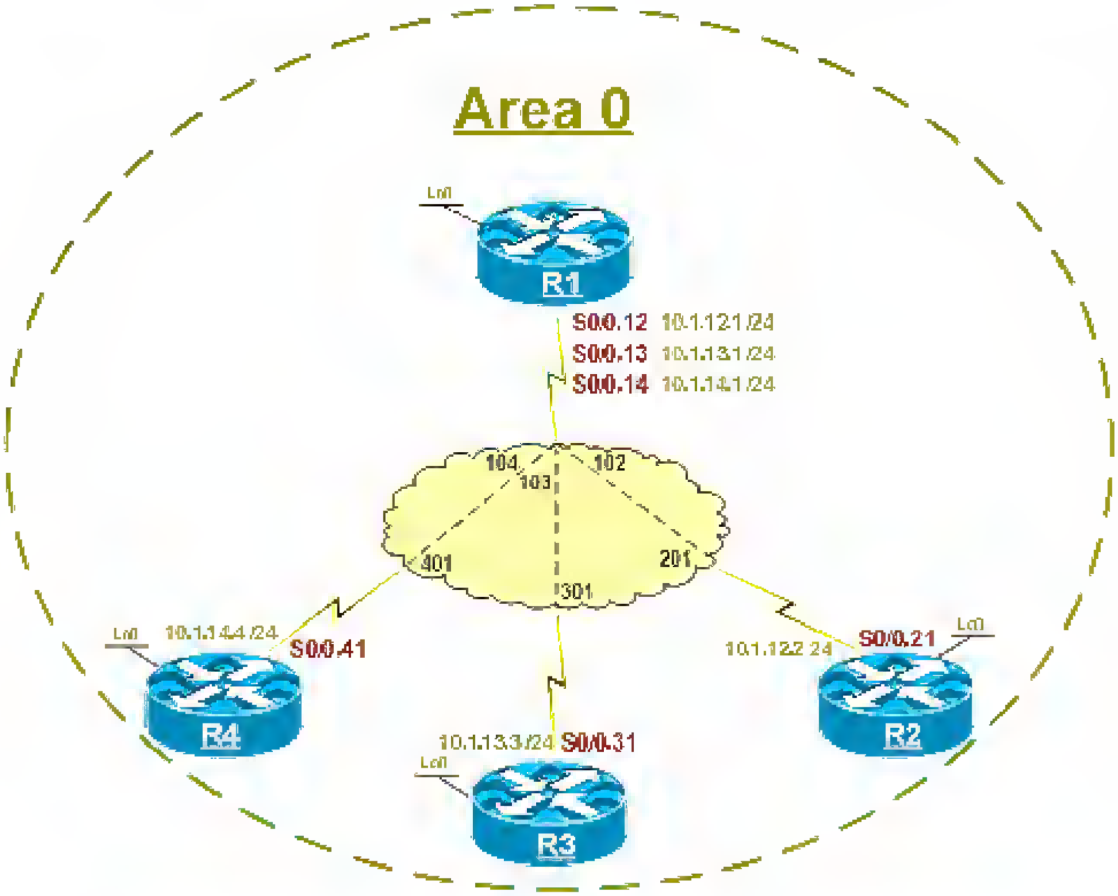

Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a Frame-relay Point-to-Point manner.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address	DLCI / Router
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24 S0/0.14 = 10.1.14.1 /24 Loopback0 = 1.1.1.1 /8	102 / R2 103 / R3 104 / R4
R2	S0/0.21 = 10.1.12.2 /24 Loopback0 = 2.2.2.2 /8	201 / R1
R3	S0/0.31 = 10.1.13.3 /24 Loopback0 = 3.3.3.3 /8	301 / R1
R4	S0/0.41 = 10.1.14.4 /24 Loopback0 = 4.4.4.4 /8	401 / R1

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in Area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. You should use the OSPF "Point-to-Point" network type to accomplish this task.

On R1

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.12.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.13.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.14.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 area 0

R2(config-router)#int lo0

R2(config-if)#ip ospf network point-to-point
```

On R3

R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.13.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point

On R4

R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.14.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point

To verify the configuration:

On R1

R1#Show ip route ospf

O   2.0.0.0/8 [110/65] via 10.1.12.2, 00:02:33, Serial0/0.12
O   3.0.0.0/8 [110/65] via 10.1.13.3, 00:02:33, Serial0/0.13
O   4.0.0.0/8 [110/65] via 10.1.14.4, 00:02:33, Serial0/0.14

On R2

R2#Show ip route ospf

O   1.0.0.0/8 [110/65] via 10.1.12.1, 00:03:07, Serial0/0.21
O   3.0.0.0/8 [110/129] via 10.1.12.1, 00:03:07, Serial0/0.21
O   4.0.0.0/8 [110/129] via 10.1.12.1, 00:03:07, Serial0/0.21
     10.0.0.0/24 is subnetted, 3 subnets
O   10.1.14.0 [110/128] via 10.1.12.1, 00:03:07, Serial0/0.21
O   10.1.13.0 [110/128] via 10.1.12.1, 00:03:07, Serial0/0.21

On R3

R3#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.13.1, 00:04:15, Serial0/0.31
O    2.0.0.0/8 [110/129] via 10.1.13.1, 00:04:15, Serial0/0.31
O    4.0.0.0/8 [110/129] via 10.1.13.1, 00:04:15, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.14.0 [110/128] via 10.1.13.1, 00:04:15, Serial0/0.31
O       10.1.12.0 [110/128] via 10.1.13.1, 00:04:15, Serial0/0.31
```

On R4

R4#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.14.1, 00:05:04, Serial0/0.41
O    2.0.0.0/8 [110/129] via 10.1.14.1, 00:05:04, Serial0/0.41
O    3.0.0.0/8 [110/129] via 10.1.14.1, 00:05:04, Serial0/0.41
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.13.0 [110/128] via 10.1.14.1, 00:05:04, Serial0/0.41
O       10.1.12.0 [110/128] via 10.1.14.1, 00:05:04, Serial0/0.41
```

Note the next hop is changed, this is because of OSPF network type, in OSPF Point-to-Point network type, the next hop IP address is no longer the router that originated the route, it's the router that advertised the route.

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 15 – OSPF Point-to-Multipoint Networks-I

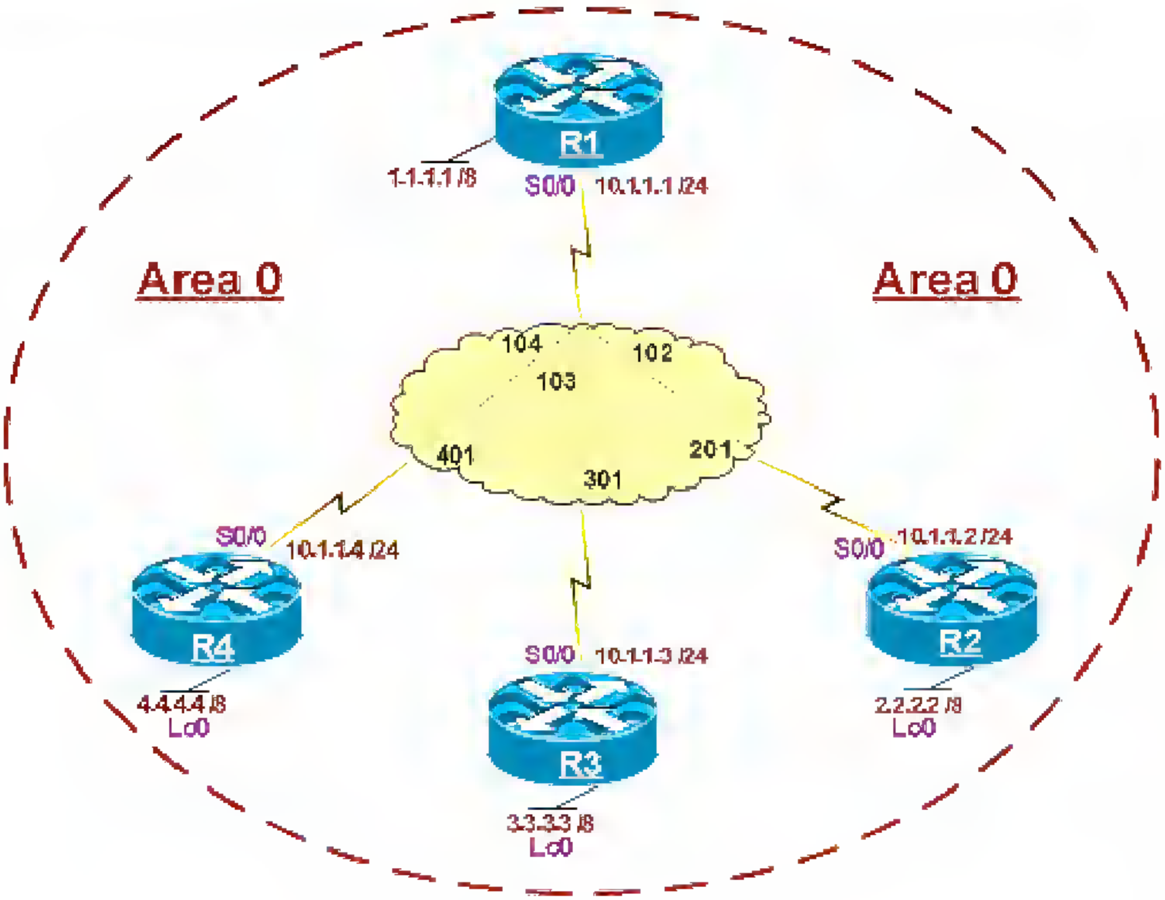

Lab Setup:

- Configure R1 as the hub and R2, R3 and R4 as spokes.
- Configure all routers in a Frame-relay Multipoint manner. DO NOT configure sub-interface/s on any of the routers. Use the "broadcast" keyword when configuring the "Frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	S0/0 = 10.1.1.2 /24 Loopback0 = 2.2.2.2 /8
R3	S0/0 = 10.1.1.3 /24 Loopback0 = 3.3.3.3 /8
R4	S0/0 = 10.1.1.4 /24 Loopback0 = 4.4.4.4 /8

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in Area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. You should use OSPF BROADCAST network type to accomplish this task.

On R1

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point

R1(config-router)#int S0/0
R1(config-if)#ip ospf network Broadcast
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```

```
R2(config-router)#int S0/0
R2(config-if)#ip ospf network Broadcast
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point

R3(config-router)#int S0/0
R3(config-if)#ip ospf network Broadcast
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point

R4(config-router)#int S0/0
R4(config-if)#ip ospf network Broadcast
```

To verify the configuration:

On R1

R1#Sh ip route ospf

```
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:46, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:46, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:00:46, Serial0/0
```

On R2

R2#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:07, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:01:07, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:01:07, Serial0/0
```

On R3

R3#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:15, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:01:15, Serial0/0
O    4.0.0.0/8 [110/65] via 10.1.1.4, 00:01:15, Serial0/0
```

On R4

R4#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:22, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:01:22, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:01:22, Serial0/0
```

Task 2

Ensure that the routers have reachability to every Loopback interface advertised in OSPF routing protocol, DO NOT use the "Frame-relay map" command or any global configuration command as part of the solution to accomplish this task.

On All Routers:

```
(config)#int s0/0
(config-if)#ip ospf net point-to-multipoint
```

To verify the configuration:

On R2

R2#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:51, Serial0/0
O    3.0.0.0/8 [110/129] via 10.1.1.1, 00:00:51, Serial0/0
O    4.0.0.0/8 [110/129] via 10.1.1.1, 00:00:51, Serial0/0
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O       10.1.1.3/32 [110/128] via 10.1.1.1, 00:00:51, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:00:51, Serial0/0
O       10.1.1.4/32 [110/128] via 10.1.1.1, 00:00:51, Serial0/0
```

On R3

R3#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:57, Serial0/0
O    2.0.0.0/8 [110/129] via 10.1.1.1, 00:00:57, Serial0/0
O    4.0.0.0/8 [110/129] via 10.1.1.1, 00:00:57, Serial0/0
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O       10.1.1.2/32 [110/128] via 10.1.1.1, 00:00:57, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:00:57, Serial0/0
O       10.1.1.4/32 [110/128] via 10.1.1.1, 00:00:57, Serial0/0
```

On R4

R4#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:06, Serial0/0
O    2.0.0.0/8 [110/129] via 10.1.1.1, 00:01:06, Serial0/0
O    3.0.0.0/8 [110/129] via 10.1.1.1, 00:01:06, Serial0/0
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O       10.1.1.2/32 [110/128] via 10.1.1.1, 00:01:06, Serial0/0
O       10.1.1.3/32 [110/128] via 10.1.1.1, 00:01:06, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:01:06, Serial0/0
```

Note OSPF Point-to-Multipoint network type creates a host route for the IP address of all the interfaces connected to the frame-relay cloud, and because of this behavior, the spoke routers can now have NLRI to all the other spoke routers, and the next hop IP address of the advertised prefixes is set based on the advertising router and NOT the router that originated the route, unless the advertising and the originator of the route happens to be the same router.

To Test the configuration:

On R2

R2#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms
```

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/124 ms

On R3

R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/120 ms

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms

On R4

R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/117 ms

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms

Task 3

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 16 – OSPF Point-to-Multipoint Networks-II

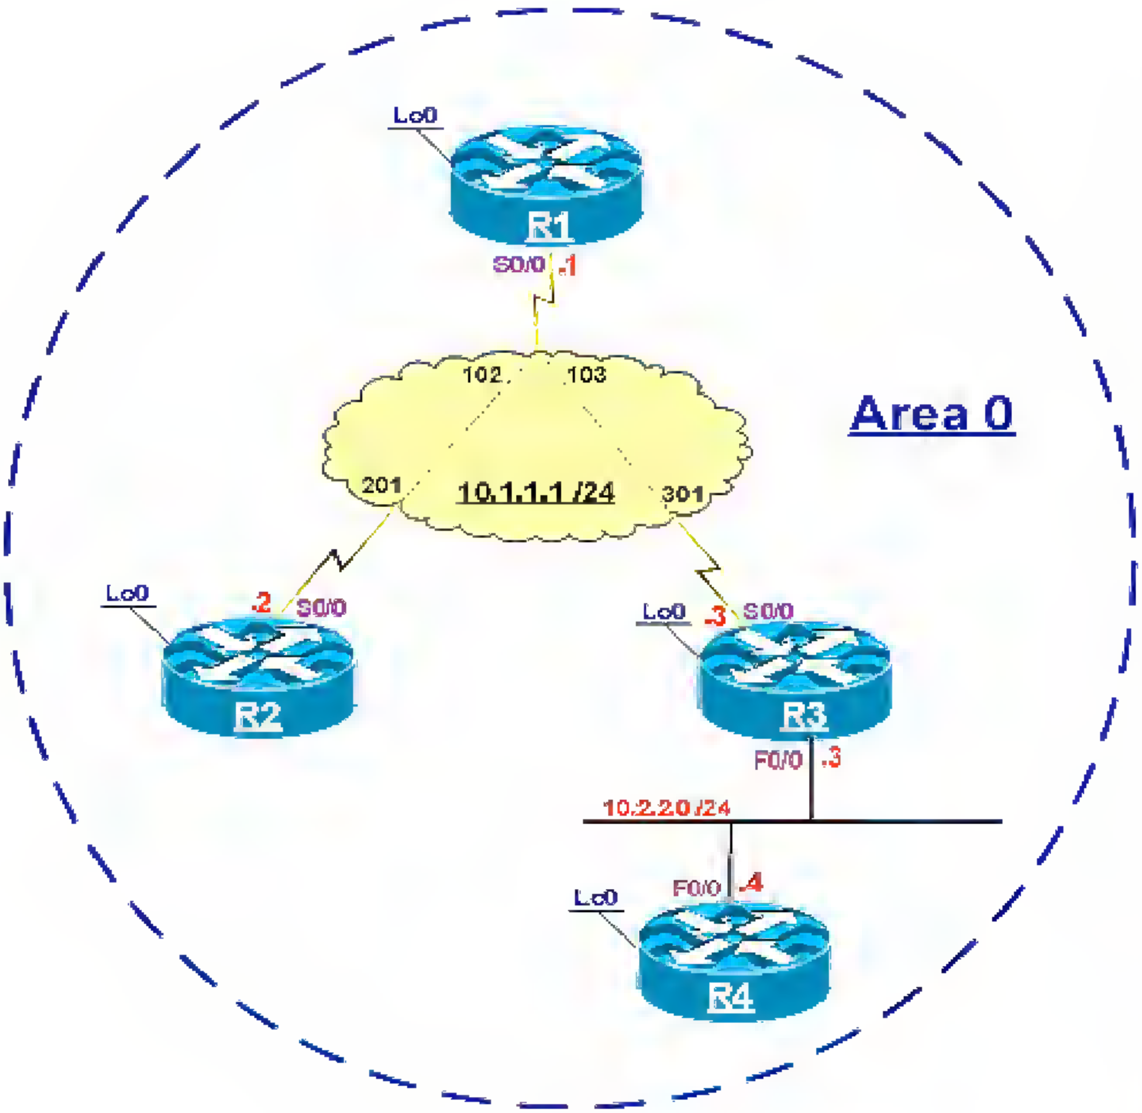

Lab Setup:

- Configure R1 as the hub and R2, R3 as spokes.
- Configure all routers in a Frame-relay Multipoint manner. DO NOT configure sub-interface/s.
- F0/0 interface of R3 and R4 should be configured in VLAN 34.
- These routers should use the "broadcast" keyword when configuring the "Frame-relay map" statements.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	S0/0 = 10.1.1.2 /24 Loopback0 = 2.2.2.2 /8
R3	S0/0 = 10.1.1.3 /24 F0/0 = 10.2.2.3 /24 Loopback0 = 3.3.3.3 /8
R4	F0/0 = 10.2.2.4 /24 Loopback0 = 4.4.4.4 /8

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in Area 0. Ensure that loopback 0 interface of these routers are advertised with their correct mask. Ensure that the OSPF BROADCAST network type is configured on the OSPF enabled interfaces except the loopbacks.

On All Routers

```
(config-if)#int lo0
(config-if)#ip ospf net point-to-point
```

On R2 and R3

(config)# int s0/0
(config-if)#ip ospf priority 0
(config-if)#ip ospf net broadcast

On R1

R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0

R1(config)#int s0/0
R1(config-if)#ip ospf priority 255

On R2

R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.1.2 0.0.0.0 are 0

On R3

R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 0
R3(config-router)#netw 10.2.2.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.1.3 0.0.0.0 are 0

On R4

R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 0
R4(config-router)#netw 10.2.2.4 0.0.0.0 are 0

To verify the configuration:

On R1

R1#Sh ip route ospf

O   2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:05, Serial0/0
O   3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:05, Serial0/0
     10.0.0.0/24 is subnetted, 2 subnets

```
O       10.2.2.0 [110/65] via 10.1.1.3, 00:00:05, Serial0/0
```

On R2

R2#Sh ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:34, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:00:34, Serial0/0
O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:00:34, Serial0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.2.2.0 [110/65] via 10.1.1.3, 00:00:34, Serial0/0
```

On R3

R3#Sh ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:05, Serial0/0
O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:00:05, Serial0/0
O    4.0.0.0/8 [110/2] via 10.2.2.4, 00:00:05, FastEthernet0/0
```

On R4

R4#Sh ip route ospf

```
O    1.0.0.0/8 [110/66] via 10.2.2.3, 00:13:48, FastEthernet0/0
O    2.0.0.0/8 [110/66] via 10.2.2.3, 00:13:48, FastEthernet0/0
O    3.0.0.0/8 [110/2] via 10.2.2.3, 00:13:48, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.1.0 [110/65] via 10.2.2.3, 00:13:48, FastEthernet0/0
```

Task 2

Ensure that these routers can Ping every loopback interface advertised in this routing domain. DO NOT use Frame-relay map, static routes, run PPP on the interfaces or any global configuration command as part of the solution to accomplish this task.

On R1, R2 and R3

```
(config)#int S0/0
(config-if)#ip ospf network point-to-multipoint
```

On R3 and R4

```
(config)#int F0/0
(config-if)#ip ospf network point-to-multipoint
```

To verify the configuration:

On R1

```
R1#Sh ip route ospf

O    2.0.0.0/8 [110/65] via 10.1.1.2, 00:02:07, Serial0/0
O    3.0.0.0/8 [110/65] via 10.1.1.3, 00:02:07, Serial0/0
O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:02:07, Serial0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.1.1.2/32 [110/64] via 10.1.1.2, 00:02:07, Serial0/0
O       10.2.2.3/32 [110/64] via 10.1.1.3, 00:02:07, Serial0/0
O       10.1.1.3/32 [110/64] via 10.1.1.3, 00:02:07, Serial0/0
O       10.2.2.4/32 [110/65] via 10.1.1.3, 00:02:07, Serial0/0
```

On R2

```
R2#Sh ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:40, Serial0/0
O    3.0.0.0/8 [110/129] via 10.1.1.1, 00:01:40, Serial0/0
O    4.0.0.0/8 [110/130] via 10.1.1.1, 00:01:40, Serial0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.2.2.3/32 [110/128] via 10.1.1.1, 00:01:40, Serial0/0
O       10.1.1.3/32 [110/128] via 10.1.1.1, 00:01:40, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:01:40, Serial0/0
O       10.2.2.4/32 [110/129] via 10.1.1.1, 00:01:40, Serial0/0
```

On R3

```
R3#Sh ip route ospf

O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:19, Serial0/0
O    2.0.0.0/8 [110/129] via 10.1.1.1, 00:01:19, Serial0/0
O    4.0.0.0/8 [110/2] via 10.2.2.4, 00:01:19, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.1.1.2/32 [110/128] via 10.1.1.1, 00:01:19, Serial0/0
O       10.1.1.1/32 [110/64] via 10.1.1.1, 00:01:19, Serial0/0
O       10.2.2.4/32 [110/1] via 10.2.2.4, 00:01:19, FastEthernet0/0
```

On R4

R4#Sh ip route ospf

O   1.0.0.0/8 [110/66] via 10.2.2.3, 00:00:40, FastEthernet0/0
O   2.0.0.0/8 [110/130] via 10.2.2.3, 00:00:40, FastEthernet0/0
O   3.0.0.0/8 [110/2] via 10.2.2.3, 00:00:40, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O   10.1.1.2/32 [110/129] via 10.2.2.3, 00:00:40, FastEthernet0/0
O   10.2.2.3/32 [110/1] via 10.2.2.3, 00:00:40, FastEthernet0/0
O   10.1.1.3/32 [110/1] via 10.2.2.3, 00:00:40, FastEthernet0/0
O   10.1.1.1/32 [110/65] via 10.2.2.3, 00:00:40, FastEthernet0/0

To test the configuration:

On R2

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/116/124 ms

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/116/124 ms

On R3

R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/124 ms

On R4

R4#Ping 1.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/61 ms

R4#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```

Task 3

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 17 – OSPF Point-to-Multipoint NON-BROADCAST Networks

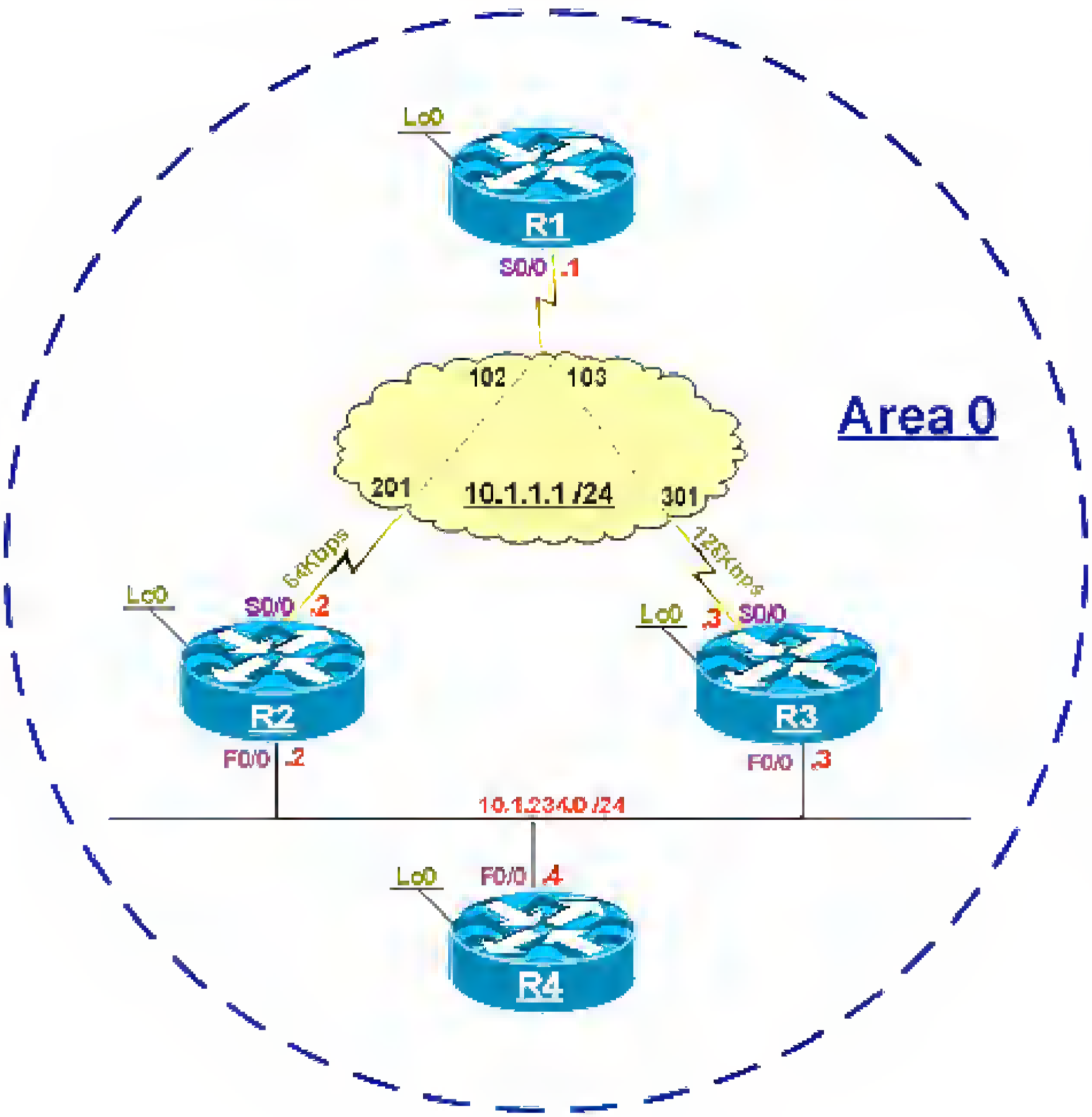

Lab Setup:

- Configure R1 as the hub and R2, and R3 as spokes.
- Configure all routers in a Frame-relay Multipoint manner. DO NOT configure sub-interfaces on any of the routers. Use the "broadcast" keyword when configuring the "frame-relay map" statements.
- Configure the F0/0 interface of R2, R3 and R4 in VLAN 234
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	S0/0 = 10.1.1.2 /24 F0/0 = 10.1.234.2 /24
R3	S0/0 = 10.1.1.3 /24 F0/0 = 10.1.234.3 /24
R4	Loopback0 = 4.4.4.4 /8 F0/0 = 10.1.234.4 /24

Task 1

Configure OSPF on all routers and advertise their directly connected interfaces in Area 0. You should use OSPF BROADCAST network type on the frame-relay interfaces to accomplish this task.
Ensure that loopback 0 interface of R1 and R4 are advertised with their correct mask.

On R1

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
```

```
R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point

R1(config-router)#int S0/0
R1(config-if)#ip ospf network Broadcast
R1(config-if)#ip ospf Priority 255
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.234.2 0.0.0.0 area 0

R2(config-router)#int S0/0
R2(config-if)#ip ospf network Broadcast
R2(config-if)#ip ospf Priority 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.234.3 0.0.0.0 area 0

R3(config-router)#int S0/0
R3(config-if)#ip ospf network Broadcast
R3(config-if)#ip ospf Priority 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.234.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1

R1#sh ip route ospf

```
O    4.0.0.0/8 [110/66] via 10.1.1.3, 00:00:21, Serial0/0
               [110/66] via 10.1.1.2, 00:00:21, Serial0/0
```

```
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.234.0 [110/65] via 10.1.1.3, 00:00:21, Serial0/0
                   [110/65] via 10.1.1.2, 00:00:21, Serial0/0
```

On R2

R2#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:01:52, Serial0/0
O    4.0.0.0/8 [110/2] via 10.1.234.4, 00:01:52, FastEthernet0/0
```

On R3

R3#Show ip route ospf

```
O    1.0.0.0/8 [110/65] via 10.1.1.1, 00:00:08, Serial0/0
O    4.0.0.0/8 [110/2] via 10.1.234.4, 00:00:08, FastEthernet0/0
```

On R4

R4#Show ip route ospf

```
O    1.0.0.0/8 [110/66] via 10.1.234.3, 00:00:19, FastEthernet0/0
               [110/66] via 10.1.234.2, 00:00:19, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.1.0 [110/65] via 10.1.234.3, 00:00:19, FastEthernet0/0
                 [110/65] via 10.1.234.2, 00:00:19, FastEthernet0/0
```

Task 2

R2 has a frame-relay cir of 64Kbps and R3 has a frame-relay cir of 128Kbps, ensure that R1 traverses through R2 to get to the networks down stream to R2 and R3, R1 should ONLY traverse through R3 if R2 is down. DO NOT use PBR to accomplish this task.

Note both R2 and R3 are advertising a cost of 1 (Ref = 100,000,000 bps / Bandwidth = 100,000,000 bps) for network 4.0.0.0 /8, R1 adds its cost of 64 through the Frame-relay interface (Ref = 100,000,000 / Bandwidth = 1,544,000 bps) /to the cost that is advertised to it by these two routers, as a result of that, R1 performs equal cost load balancing, remember R2 or R3's frame-relay cost is NOT calculated.

One possible method of dealing with this scenario is to configure R1's frame-relay

interface with OSPF Point-to-Multipoint Non-Broadcast network type, this network type allows R1 to associate a cost to each of it's downstream neighbors, the neighbor with a lower cost will be chosen as the best route.
Remember in order for 2 OSPF routers to exchange routes the network types must match, but there are 2 exceptions to this rule and they are as follows:

1. A Point-to-Multipoint ←→ Point-to-Point
2. A Broadcast ←→ Non-broadcast

In this case the first option is exercised as follows:

On R1

```
R1(config)#int S0/0
R1(config-if)#ip ospf network point-to-multipoint non-broadcast
```

The following command changes the OSPF hello-interval to match R2 and R3's hello-interval

```
R1(config-if)#ip ospf hello-interval 10
```

```
R1(config-if)#router ospf 1
R1(config-router)#neighbor 10.1.1.2 cost 1
R1(config-router)#neighbor 10.1.1.3 cost 2
```

The following command changes the network type of R2 and R3 to point-to-point.

On R2 and R3

```
(config)#int S0/0
(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1

R1#Show ip route ospf

```
O    4.0.0.0/8 [110/3] via 10.1.1.2, 00:00:21, Serial0/0
     10.0.0.0/24 is subnetted, 2 subnets
O       10.1.234.0 [110/2] via 10.1.1.2, 00:00:21, Serial0/0
```

Task 3

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 18
OSPF and NBMA

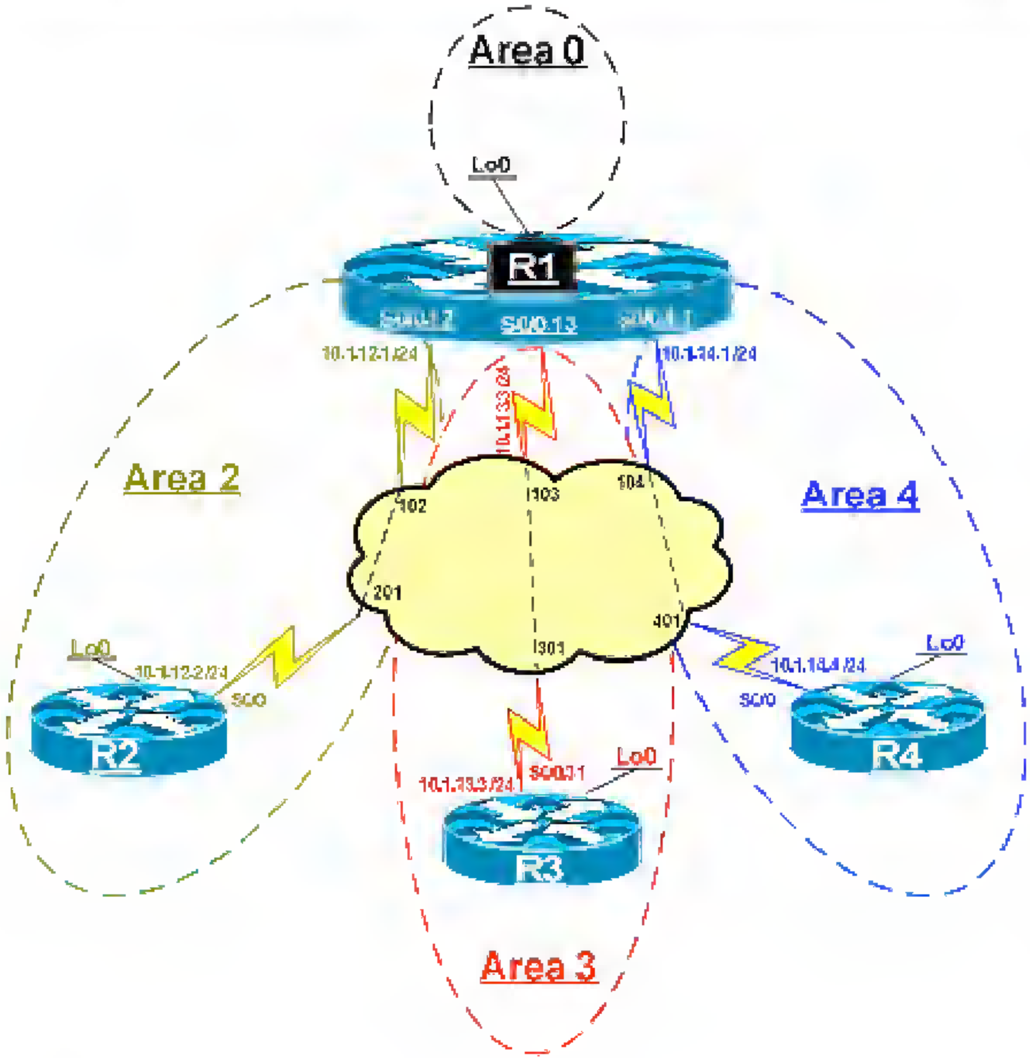

Lab Setup:

- R1 should be configured with three sub-interfaces; the first sub-interface should be configured in a point-to-point manner connecting R1 to R2. R2 should not use a sub-interface for this connection.
- The second and the third sub-interface of R1 should be configured in a multipoint manner, one connecting R1 to R3 and the other one connecting R1 to R4.
- R3 should be configured in a point-to-point manner.
- R4 should NOT use a sub-interface for it's connection to R1.

IP Addressing:

Router	Interface	IP address	Area
R1	Lo0 S0/0.12 S0/0.13 S0/0.14	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24 10.1.14.1 /24	Area 0 Area 2 Area 3 Area 4
R2	Lo0 S0/0	2.2.2.2 /8 10.1.12.2 /24	Area 2 Area 2
R3	Lo0 S0/0.31	3.3.3.3 /8 10.1.13.3 /24	Area 3 Area 3
R4	Lo0 S0/0	4.4.4.4 /8 10.1.14.4 /24	Area 4 Area 4

R1 should be the DR in all cases, if one is required.

Task 1

Configure OSPF on all routers and advertise their directly connected networks in their assigned area identified in the IP addressing chart.

On R1

```
R1(config)#router ospf 1
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 2
R1(config-router)#netw 10.1.13.1 0.0.0.0 are 3
R1(config-router)#netw 10.1.14.1 0.0.0.0 are 4
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 0.0.0.0 0.0.0.0 area 2
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 0.0.0.0 0.0.0.0 area 3
```

On R4

```
R4(config-if)#router ospf 1
R4(config-router)#netw 0.0.0.0 0.0.0.0 area 4
```

Task 2

Ensure that when the routers in area 2 attempt to establish a neighbor adjacency they are successful (FULL STATE), but no routes are exchanged. DO NOT configure R2 to accomplish this task.

On R1

```
R1(config)#int S0/0.12
R1(config-subif)#ip ospf network point-to-multipoint non-broadcast

R1(config)#router ospf 1
R1(config-router)#neighbor 10.1.12.2
```

To verify the configuration:

On R1

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/  -        00:01:49    10.1.12.2       Serial0/0.12
3.3.3.3           0   FULL/  -        00:00:38    10.1.13.3       Serial0/0.13

R1#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B – BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route

       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C       1.0.0.0/8 is directly connected, Loopback0
```

```
C       10.1.14.0 is directly connected, Serial0/0.14
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

Note the two routers are in Full state but they have not exchanged routes. This can also be accomplished with "point-to-multipoint" network type.

Task 3

Ensure that the routers in area 2 can establish an OSPF neighbor adjacency. R2 should not be configured at all to accomplish this task.

On R1

```
R1(config)#int S0/0.12
R1(config-subif)#ip ospf network non-broadcast
R1(config-subif)#ip ospf priority 255

R1(config)#router ospf 1
R1(config-router)#neighbor 10.1.12.2
```

Note the "ip ospf priority" command is required to make R1 the DR. When frame-relay is configured directly under the physical interface (Multipoint), the OSPF network type will default to non-broadcast. In this task the network type of R1's interface s0/0.12 is also changed to non-broadcast to match R2's network type. In non-broadcast networks the "neighbor" command in router config mode must be configured so the OSPF hello packets are exchanged via Unicast.

To Verify the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B – BGP

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:08, Serial0/0.12
     10.1.0.0/24 is subnetted, 3 subnets
C       10.1.14.0 is directly connected, Serial0/0.14
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

Task 4

Area 3 should be configured in a point-to-point network type, only one of the routers should be changed to accomplish this task.

On R1

```
R1(config)#int S0/0.13
R1(config-subif)#ip ospf net point-to-point
```

In the earlier IOS releases when an interface was changed from "non-broadcast" to "point-to-point" we had to change the hello interval as well, because if the hello intervals did not match, the routers did not form neighbor adjacency. In the latest IOS releases the hello intervals automatically change when the network type is changed.

To verify the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B – BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route
```

```
Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:04:11, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/65] via 10.1.13.3, 00:01:26, Serial0/0.13
     10.1.0.0/24 is subnetted, 3 subnets
C       10.1.14.0 is directly connected, Serial0/0.14
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

Task 5

Area 4 should be configured with a totally different network type than task 2, 3 and 4. DO NOT use point-to-multipoint to accomplish this task.

On R1

```
R1(config)#int S0/0.14
R1(config-subif)#ip ospf net broadcast
R1(config-subif)#ip ospf priority 255
```

On R4

```
R4(config-if)#int S0/0
R4(config-if)#ip ospf net broadcast
```

To verify the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 10.1.12.2, 00:10:54, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/65] via 10.1.13.3, 00:08:08, Serial0/0.13
     4.0.0.0/32 is subnetted, 1 subnets
O       4.4.4.4 [110/65] via 10.1.14.4, 00:00:01, Serial0/0.14
     10.1.0.0/24 is subnetted, 3 subnets
C       10.1.14.0 is directly connected, Serial0/0.14
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

Task 6

Remove the priority command from R1's S0/0.12 and set the network type to "point-to-multipoint non-broadcast". Ensure that these routers exchange routes. Do NOT change the network type to accomplish this task.

On R1

```
R1(config)#int S0/0.12
R1(config-subif)#NO ip ospf priority 255
R1(config-subif)#ip ospf network point-to-multipoint non-broadcast

R1(config)#Interface Tunnel 1
R1(config-if)#ip address 200.1.12.1 255.255.255.0
R1(config-if)#tunnel source 10.1.12.1
R1(config-if)#tunnel destination 10.1.12.2

R1(config-if)#router ospf 1
R1(config-router)#NO netw 10.1.12.1 0.0.0.0 area 2
R1(config-router)#netw 200.1.12.1 0.0.0.0 area 0
```

The reason to remove the network 10.1.12.0 from OSPF is to prevent recursive loops.

On R2

```
R2(config)#Interface Tunnel 1
R2(config-if)#ip address 200.1.12.2 255.255.255.0
R2(config-if)#tunnel source 10.1.12.2
R2(config-if)#tunnel destination 10.1.12.1

R2(config-if)#router ospf 1
R2(config-router)#NO netw 0.0.0.0 0.0.0.0 area 2
R2(config-router)#netw 200.1.12.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 2
```

Note on R2 we must remove the earlier network statement (0.0.0.0 0.0.0.0 area 2) or else all the interfaces will be advertised in area 2 and this is not the desired behavior. Lastly we must advertise network 2.0.0.0 in area 2.

To verify the configuration:

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

O      1.0.0.0/8 [110/11112] via 200.1.12.1, 00:00:11, Tunnel1
C      2.0.0.0/8 is directly connected, Loopback0
O IA  3.0.0.0/8 [110/11176] via 200.1.12.1, 00:00:05, Tunnel1
         4.0.0.0/32 is subnetted, 1 subnets
O IA    4.4.4.4 [110/11176] via 200.1.12.1, 00:00:11, Tunnel1
C       200.1.12.0/24 is directly connected, Tunnel1
          10.1.0.0/24 is subnetted, 3 subnets
O IA    10.1.14.0 [110/11175] via 200.1.12.1, 00:00:11, Tunnel1
C        10.1.12.0 is directly connected, Serial0/0
O IA    10.1.13.0 [110/11175] via 200.1.12.1, 00:00:12,
```

Note the reason network 1.0.0.0 shows up as an Intra-area route is because R2 has an interface in area 0 and network 1.0.0.0 is from area 0.

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C      1.0.0.0/8 is directly connected, Loopback0
O IA 2.0.0.0/8 [110/11112] via 200.1.12.2, 00:02:33, Tunnel1
O      3.0.0.0/8 [110/65] via 10.1.13.3, 00:02:43, Serial0/0.13
O      4.0.0.0/8 [110/65] via 10.1.14.4, 00:02:33, Serial0/0.14
C      200.1.12.0/24 is directly connected, Tunnel1
       10.1.0.0/24 is subnetted, 3 subnets
C       10.1.14.0 is directly connected, Serial0/0.14
C       10.1.12.0 is directly connected, Serial0/0.12
C       10.1.13.0 is directly connected, Serial0/0.13
```

The reason network 2.0.0.0 shows up as an Inter-area route is because the local router (R1) does not have an interface in area 2.

Task 7

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 19
Forward Address Suppression

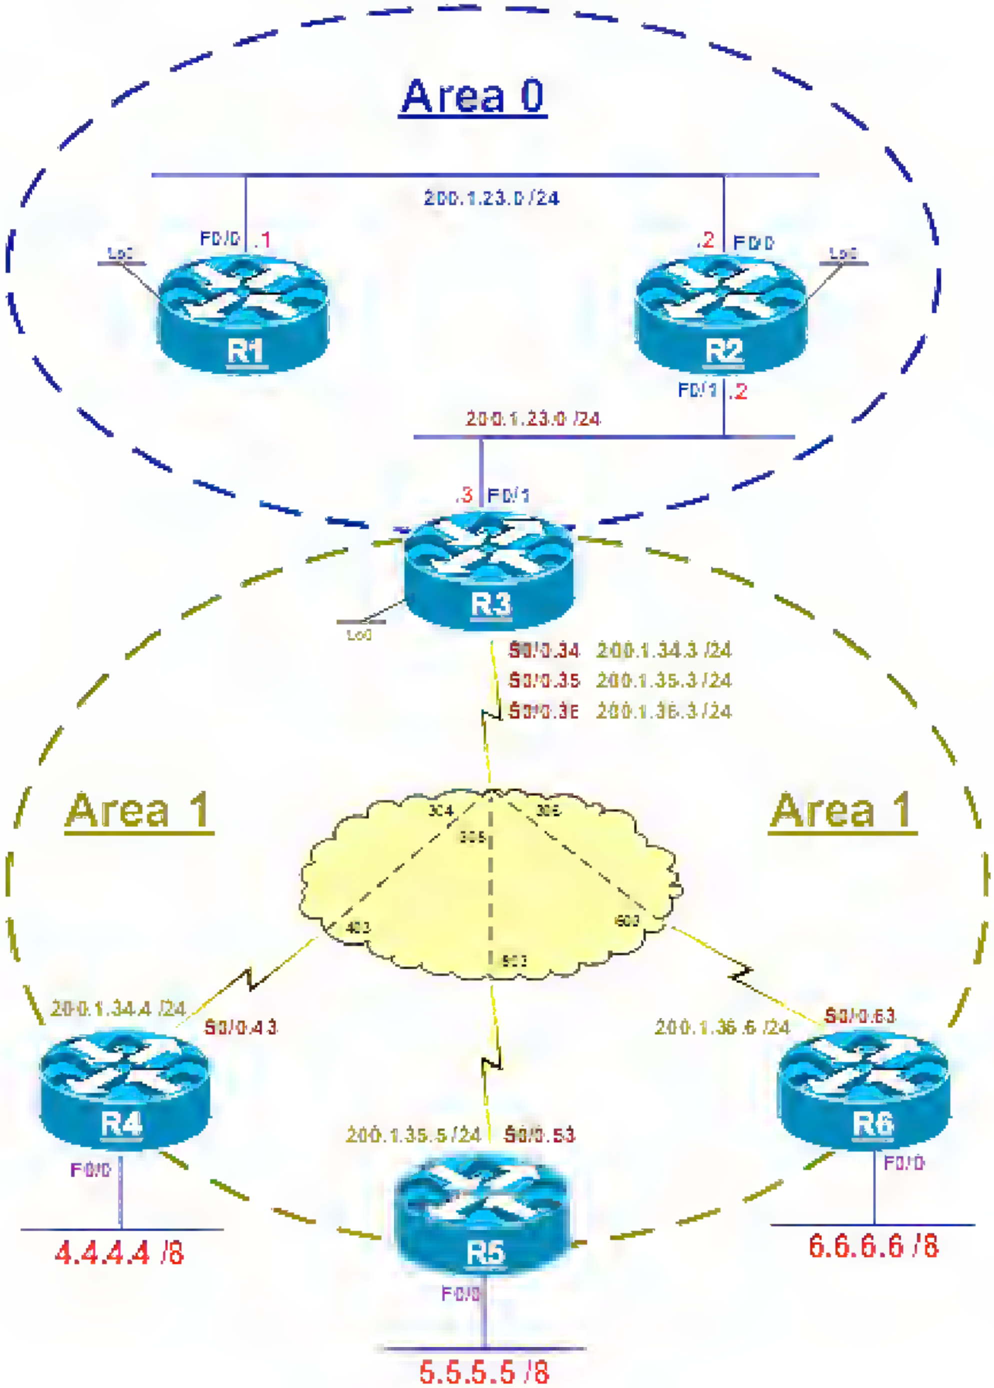

Lab Setup:

- Every frame-relay connection should be configured in a Point-to-Point manner
- Use the IP addressing scheme below for IP address assignment:
- Configure the F0/0 interface of R1 and R2 in VLAN 12 and the F0/1 interface of R2 and R3 in VLAN 23; F0/0 interface of R4, R5 and R6 should be configured in VLAN 4, 5 and 6 respectively.

IP addressing:

Router	Interface / IP address	DLCI	Connecting to
R1	S0/0.12 = 200.1.12.1 /24 Lo0 = 1.1.1.1 /8	102	R2
R2	F0/0 = 200.1.12.2 /24 F0/1 = 200.1.23.2 /24 Lo0 = 2.2.2.2 /8		R1 R3
R3	F0/0 = 200.1.23.3 /24 S0/0.34 = 200.1.34.3 /24 S0/0.35 = 200.1.35.3 /24 S0/0.36 = 200.1.36.3 /24 Lo0 = 3.3.3.3 /8	 304 305 306	R2 R4 R5 R6
R4	S0/0.43 = 200.1.34.4 /24 F0/0 = 4.4.4.4 /8	403	R3
R5	S0/0.53 = 200.1.35.5 /24 F0/0 = 5.5.5.5 /8	503	R3
R6	S0/0.63 = 200.1.36.6 /24 F0/0 = 6.6.6.6 /8	603	R3

Task 1

Configure OSPF Area 0 on R1, R2 and advertise their directly connected interfaces in this Area.

On R1 and R2

```
(config)#router ospf 1
(config-router)#netw 0.0.0.0 0.0.0.0 area 0
```

To verify the configuration:

On R1

R1#Sh ip route ospf

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:00:01, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:00:01, Serial0/0.12
```

Task 2

Configure R3's Loopback interface and its frame-relay connection to R4, R5 and R6 in Area 1, and it's Frame-relay connection to R2 in Area 0.

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.34.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.35.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.36.3 0.0.0.0 area 1
R3(config-router)#netw 200.1.23.3 0.0.0.0 area 0
```

To verify the configuration:

On R3

R3#Sh ip route ospf

```
     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/846] via 200.1.23.2, 00:00:58, Serial0/1.32
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/782] via 200.1.23.2, 00:00:58, Serial0/1.32
O    200.1.12.0/24 [110/845] via 200.1.23.2, 00:00:58, Serial0/1.32
```

On R1

R1#Sh ip route ospf

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:02:05, Serial0/0.12
```

```
O IA 200.1.36.0/24 [110/909] via 200.1.12.2, 00:02:05, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:02:05, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/129] via 200.1.12.2, 00:02:05, Serial0/0.12
O IA 200.1.34.0/24 [110/909] via 200.1.12.2, 00:02:05, Serial0/0.12
O IA 200.1.35.0/24 [110/909] via 200.1.12.2, 00:02:05, Serial0/0.12
```

Task 3

Configure the Frame-relay connection of R4, R5 and R6 to R3 in Area 1. These routers should redistribute their F0/0 interface in OSPF routing protocol, you should NOT use an access-list or a prefix-list to accomplish this task.

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 200.1.34.4 0.0.0.0 area 1
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#netw 200.1.35.5 0.0.0.0 area 1
```

On R6

```
R6(config)#router ospf 1
R6(config-router)#netw 200.1.36.6 0.0.0.0 area 1
```

On R4, R5 and R6

```
(config)#route-map TST permit 10
(config-route-map)#match interface F0/0

(config)#router ospf 1
(config-router)#redistribute connected route-map TST subnets
```

To verify the configuration:

On R1

```
R1#Sh ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:12:37, Serial0/0.12
O IA 200.1.36.0/24 [110/909] via 200.1.12.2, 00:12:37, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:12:37, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/129] via 200.1.12.2, 00:12:37, Serial0/0.12
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:03:39, Serial0/0.12
O IA 200.1.34.0/24 [110/909] via 200.1.12.2, 00:12:37, Serial0/0.12
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:00:14, Serial0/0.12
O IA 200.1.35.0/24 [110/909] via 200.1.12.2, 00:12:37, Serial0/0.12
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:01:29, Serial0/0.12
```

Task 4

Configure Area 1 as a NSSA.

On R3, R4, R5 and R6

```
(config)#router ospf 1
(config-router)#area 1 nssa
```

To verify the configuration:

On R3

```
R3#Sh ip route ospf

     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/846] via 200.1.23.2, 00:01:30, Serial0/1.32
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/782] via 200.1.23.2, 00:01:30, Serial0/1.32
O N2 4.0.0.0/8 [110/20] via 200.1.34.4, 00:00:35, Serial0/1.34
O N2 5.0.0.0/8 [110/20] via 200.1.35.5, 00:00:35, Serial0/1.35
O N2 6.0.0.0/8 [110/20] via 200.1.36.6, 00:00:35, Serial0/1.36
O    200.1.12.0/24 [110/845] via 200.1.23.2, 00:01:30, Serial0/1.32
```

On R1

```
R1#Sh ip route ospf
```

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:02:37, Serial0/0.12
O IA 200.1.36.0/24 [110/909] via 200.1.12.2, 00:02:37, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:02:37, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/129] via 200.1.12.2, 00:02:37, Serial0/0.12
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:02:01, Serial0/0.12
O IA 200.1.34.0/24 [110/909] via 200.1.12.2, 00:02:37, Serial0/0.12
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:01:51, Serial0/0.12
O IA 200.1.35.0/24 [110/909] via 200.1.12.2, 00:02:37, Serial0/0.12
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:01:41, Serial0/0.12
```

On R2

R2#Sh ip route ospf

```
     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/65] via 200.1.12.1, 00:04:01, Serial0/0.21
O IA 200.1.36.0/24 [110/845] via 200.1.23.3, 00:04:01, Serial0/0.23
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/65] via 200.1.23.3, 00:04:01, Serial0/0.23
O E2 4.0.0.0/8 [110/20] via 200.1.23.3, 00:03:25, Serial0/0.23
O IA 200.1.34.0/24 [110/845] via 200.1.23.3, 00:04:01, Serial0/0.23
O E2 5.0.0.0/8 [110/20] via 200.1.23.3, 00:03:15, Serial0/0.23
O IA 200.1.35.0/24 [110/845] via 200.1.23.3, 00:04:01, Serial0/0.23
O E2 6.0.0.0/8 [110/20] via 200.1.23.3, 00:03:05, Serial0/0.23
```

Task 5

Configure R3 to filter the following networks:
200.1.34.0 /24, 200.1.35.0 /24 and 200.1.36.0 /24

On R3

```
R3(config)#router ospf 1
R3(config-router)#area 1 range 200.1.34.0 255.255.255.0 not-advertise
R3(config-router)#area 1 range 200.1.35.0 255.255.255.0 not-advertise
R3(config-router)#area 1 range 200.1.36.0 255.255.255.0 not-advertise
```

To verify the configuration:

On R2

R2#Sh ip route ospf

```
     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/65] via 200.1.12.1, 00:01:47, Serial0/0.21
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/65] via 200.1.23.3, 00:01:47, Serial0/0.23
```

On R1

R1#Sh ip route ospf

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:03:02, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:03:02, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/129] via 200.1.12.2, 00:03:02, Serial0/0.12
```

Note the routers in Area 0 no longer have reachability to the prefixes from Area 1.

Task 6

Configure R3 such that the routers in Area 0 can reach the networks that were redistributed in step 3. Use minimum number of commands to accomplish this task; you should NOT use any global config commands as part of the solution to this task. DO NOT remove the commands from the previous step.

On R3

```
R3(config)#router ospf 1
R3(config-router)#area 1 nssa translate type7 suppress-fa
```

To verify the configuration:

On R2

R2#Sh ip route ospf

```
     1.0.0.0/32 is subnetted, 1 subnets
```

```
O       1.1.1.1 [110/65] via 200.1.12.1, 00:07:29, Serial0/0.21
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/65] via 200.1.23.3, 00:07:29, Serial0/0.23
O E2 4.0.0.0/8 [110/20] via 200.1.23.3, 00:00:41, Serial0/0.23
O E2 5.0.0.0/8 [110/20] via 200.1.23.3, 00:00:41, Serial0/0.23
O E2 6.0.0.0/8 [110/20] via 200.1.23.3, 00:00:41, Serial0/0.23
```

On R1

R1#Show ip route ospf

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:08:03, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:08:03, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/129] via 200.1.12.2, 00:08:03, Serial0/0.12
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:01:15, Serial0/0.12
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:01:15, Serial0/0.12
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:01:15, Serial0/0.12
```

To test the configuration:

On R1

R1#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

R1#Ping 4.4.4.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 168/171/173 ms
```

R1#Ping 5.5.5.5

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 168/170/173 ms
```

The OSPF Forwarding Address Suppression in Translated Type-5 LSAs feature causes an NSSA ABR to translate Type-7 LSAs to Type-5 LSAs, but use the 0.0.0.0 as the forwarding address instead of that specified in the Type-7 LSA. Note if the "Area 1 translate type7 suppress-fa" command and the filters were removed, this will be the output of the "Show ip ospf data external 4.0.0.0"

R1#Sh ip ospf data external 4.0.0.0

```
        OSPF Router with ID (1.1.1.1) (Process ID 1)

           Type-5 AS External Link States

  Routing Bit Set on this LSA
  LS age: 7
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 4.0.0.0 (External Network Number )
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000003
Checksum: 0xF5A8
  Length: 36
  Network Mask: /8
       Metric Type: 2 (Larger than any link state path)
       TOS: 0
       Metric: 20
       Forward Address: 200.1.34.4
       External Route Tag: 0
```

This is the router that advertised the network to the local router

Note the address is not suppressed. Basically the IP address of the that originated the route

After the filters are applied and the "Area 1 nssa translate type7 suppress-fa" command is configured, the output of the "Show ip ospf da ext 4.0.0.0" will be changed as follows:

R1#Sh ip ospf data external 4.0.0.0

```
        OSPF Router with ID (1.1.1.1) (Process ID 1)

           Type-5 AS External Link States


Routing Bit Set on this LSA
  LS age: 293
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 4.0.0.0 (External Network Number )
  Advertising Router: 3.3.3.3
```

Note the advertising router is still in the DB

```
LS Seq Number: 80000002
Checksum: 0x5738
Length: 36
Network Mask: /8
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 20
      Forward Address: 0.0.0.0
      External Route Tag: 0
```

The IP address of the router that originated the route is suppressed.

Because the IP address of the router that originated the route/s are suppressed, area 0 routers no longer need to maintain extra prefixes in their routing table.

On R2

R2#Sh ip route ospf

```
     1.0.0.0/32 is subnetted, 1 subnets
O       1.1.1.1 [110/65] via 200.1.12.1, 00:07:29, Serial0/0.21
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/65] via 200.1.23.3, 00:07:29, Serial0/0.23
O E2 4.0.0.0/8 [110/20] via 200.1.23.3, 00:00:41, Serial0/0.23
O E2 5.0.0.0/8 [110/20] via 200.1.23.3, 00:00:41, Serial0/0.23
O E2 6.0.0.0/8 [110/20] via 200.1.23.3, 00:00:41, Serial0/0.23
```

On R1

R1#Show ip route ospf

```
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/65] via 200.1.12.2, 00:08:03, Serial0/0.12
O    200.1.23.0/24 [110/128] via 200.1.12.2, 00:08:03, Serial0/0.12
     3.0.0.0/32 is subnetted, 1 subnets
O IA    3.3.3.3 [110/129] via 200.1.12.2, 00:08:03, Serial0/0.12
O E2 4.0.0.0/8 [110/20] via 200.1.12.2, 00:01:15, Serial0/0.12
O E2 5.0.0.0/8 [110/20] via 200.1.12.2, 00:01:15, Serial0/0.12
O E2 6.0.0.0/8 [110/20] via 200.1.12.2, 00:01:15, Serial0/0.12
```

Note the backbone routers no longer need to maintain the extra prefixes for the links, but they have full reachability to the prefixes that were redistributed.

Task 7

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 20 – OSPF NSSA no-redistribution & injection of default routes

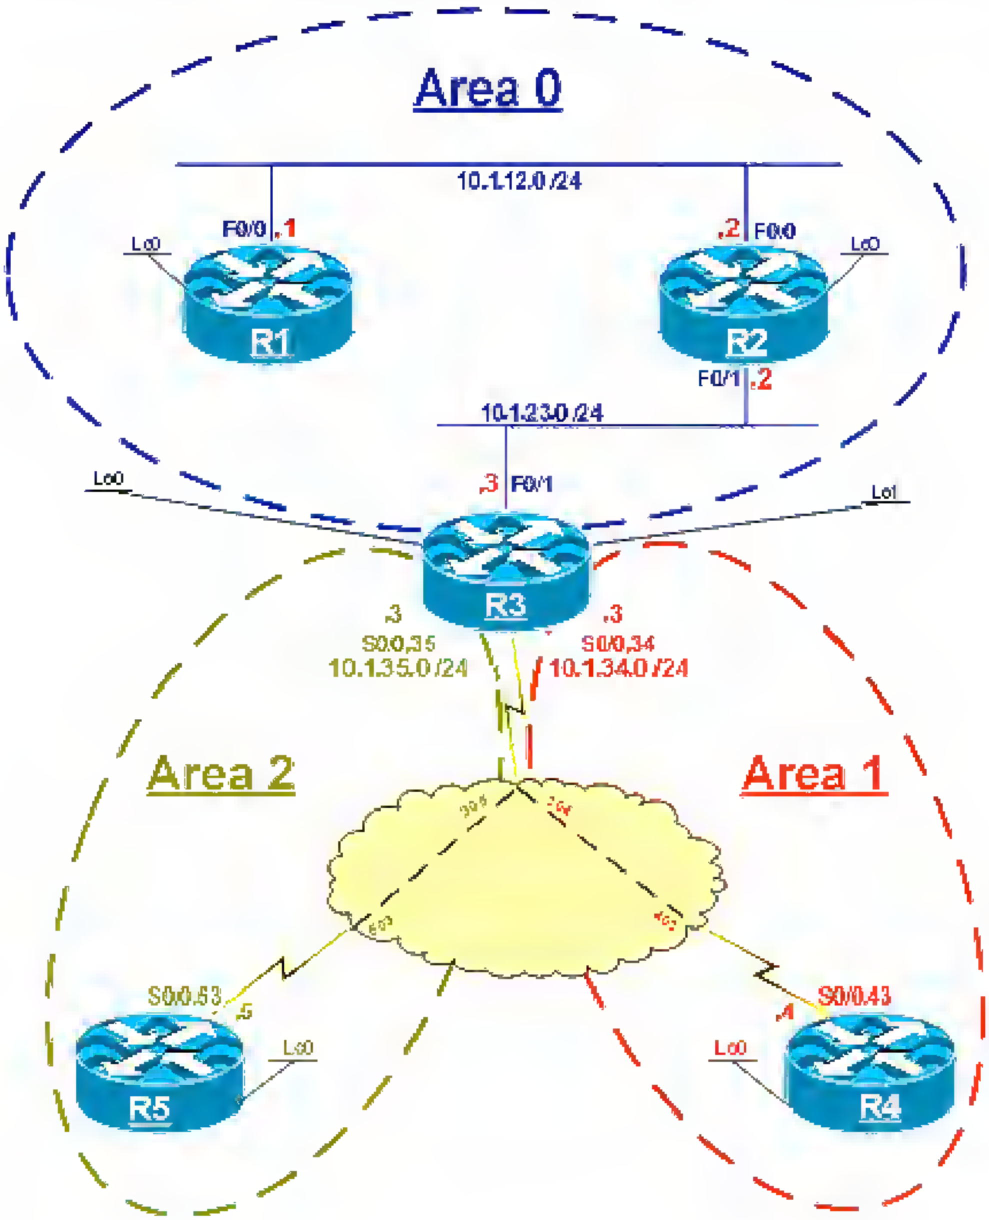

Lab Setup:

- Configure all frame-relay connections in a point-to-point manner.
- Use the IP addressing and DLCI chart below.

IP addressing:

Router	Interface / IP address	DLCI assignment
R1	F0/0 = 10.1.12.1 /24 Loopback0 = 1.1.1.1 /8	
R2	**F0/0 = 10.1.12.2 /24** **F0/1 = 10.1.23.2 /24** **Loopback0 = 2.2.2.2 /8**	
R3	F0/1 = 10.1.23.3 /24 S0/0.34 = 10.1.34.3 /24 S0/0.35 = 10.1.35.3 /24 Loopback0 = 3.3.3.3 /8 Loopback1 = 33.3.3.3 /8	 304 305
R4	**S0/0.43 = 10.1.34.4 /24** **Loopback0 = 4.4.4.4 /8**	403
R5	S0/0.53 = 10.1.35.5 /24 Loopback0 = 5.5.5.5 /8	503

Task 1

Configure OSPF on the routers based on the following chart:

Router	Interface / Area
R1	S0/0.12 / **Area 0** Loopback0 / **Area 0**
R2	S0/0.21 / **Area 0** S0/0.23 / **Area 0** Loopback0 / **Area 0**
R3	S0/0.32 / **Area 0** S0/0.34 / **Area 1** S0/0.35 / **Area 2**
R4	S0/0.43 / **Area 1** Loopback0 / **Area 1**
R5	S0/0.53 / **Area 2** Loopback0 / **Area 2**

On R1

```
R1(config-if)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 0
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.23.2 0.0.0.0 are 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 1
R3(config-router)#netw 10.1.35.3 0.0.0.0 area 2
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 1
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 2
R5(config-router)#netw 10.1.35.5 0.0.0.0 area 2
```

To verify the configuration:

On R1

R1#Show ip route ospf | Inc O

```
O       2.2.2.2 [110/65] via 10.1.12.2, 00:05:31, Serial0/0.12
O IA    4.4.4.4 [110/193] via 10.1.12.2, 00:03:49, Serial0/0.12
O IA    5.5.5.5 [110/193] via 10.1.12.2, 00:02:31, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:05:31, Serial0/0.12
O IA    10.1.35.0 [110/192] via 10.1.12.2, 00:05:21, Serial0/0.12
```

```
O IA    10.1.34.0 [110/192] via 10.1.12.2, 00:05:31, Serial0/0.12
```

On R2

R2#Show ip route ospf | Inc O

```
O       1.1.1.1 [110/65] via 10.1.12.1, 00:06:22, Serial0/0.21
O IA    4.4.4.4 [110/129] via 10.1.23.3, 00:04:40, Serial0/0.23
O IA    5.5.5.5 [110/129] via 10.1.23.3, 00:03:22, Serial0/0.23
O IA    10.1.35.0 [110/128] via 10.1.23.3, 00:06:12, Serial0/0.23
O IA    10.1.34.0 [110/128] via 10.1.23.3, 00:06:22, Serial0/0.23
```

On R3

R3#Show ip route ospf | Inc O

```
O       1.1.1.1 [110/129] via 10.1.23.2, 00:07:00, Serial0/0.32
O       2.2.2.2 [110/65] via 10.1.23.2, 00:07:00, Serial0/0.32
O       4.4.4.4 [110/65] via 10.1.34.4, 00:05:28, Serial0/0.34
O       5.5.5.5 [110/65] via 10.1.35.5, 00:04:10, Serial0/0.35
O       10.1.12.0 [110/128] via 10.1.23.2, 00:07:00, Serial0/0.32
```

On R4

R4#Show ip route ospf | Inc O

```
O IA    1.1.1.1 [110/193] via 10.1.34.3, 00:06:18, Serial0/0.43
O IA    2.2.2.2 [110/129] via 10.1.34.3, 00:06:18, Serial0/0.43
O IA    5.5.5.5 [110/129] via 10.1.34.3, 00:05:00, Serial0/0.43
O IA    10.1.12.0 [110/192] via 10.1.34.3, 00:06:18, Serial0/0.43
O IA    10.1.23.0 [110/128] via 10.1.34.3, 00:06:18, Serial0/0.43
O IA    10.1.35.0 [110/128] via 10.1.34.3, 00:06:18, Serial0/0.43
```

On R5

R5#Show ip route ospf | Inc O

```
O IA    1.1.1.1 [110/193] via 10.1.35.3, 00:06:02, Serial0/0.53
O IA    2.2.2.2 [110/129] via 10.1.35.3, 00:06:02, Serial0/0.53
O IA    4.4.4.4 [110/129] via 10.1.35.3, 00:06:02, Serial0/0.53
O IA    10.1.12.0 [110/192] via 10.1.35.3, 00:06:02, Serial0/0.53
O IA    10.1.23.0 [110/128] via 10.1.35.3, 00:06:02, Serial0/0.53
O IA    10.1.34.0 [110/128] via 10.1.35.3, 00:06:02, Serial0/0.53
```

Task 2

Configure R3 to redistribute its Loopback 0 and 1 interfaces into this OSPF routing domain.

On R3

```
R3(config)#route-map TST permit 10
R3(config-route-map)#match interface lo0 lo1

R3(config)#router ospf 1
R3(config-router)#redistribute connected subnets route-map TST
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc E2

O E2 33.0.0.0/8 [110/20] via 10.1.12.2, 00:01:10, Serial0/0.12
O E2 3.0.0.0/8 [110/20] via 10.1.12.2, 00:01:10, Serial0/0.12
```

On R5

```
R5#Show ip route ospf | Inc E2

O E2 33.0.0.0/8 [110/20] via 10.1.35.3, 00:02:09, Serial0/0.53
O E2 3.0.0.0/8 [110/20] via 10.1.35.3, 00:02:09, Serial0/0.53
```

Task 3

Configure area 1 and area 2 as NSSA, R3 should be configured such that the routers in these two areas get a default route, this default route should be injected as an external route.

On R3

```
R3(config)#Router ospf 1
R3(config-router)#area 1 nssa default-information-originate
R3(config-router)#area 2 nssa default-information-originate
```

On R4

R4(config)#router ospf 1
R4(config-router)#area 1 nssa

On R5

R5(config)#router ospf 1
R5(config-router)#area 2 nssa

To verify the configuration:

On R4

R4#Show ip route ospf | Inc O

```
O IA    1.1.1.1 [110/193] via 10.1.34.3, 00:05:05, Serial0/0.43
O IA    2.2.2.2 [110/129] via 10.1.34.3, 00:05:05, Serial0/0.43
O N2 33.0.0.0/8 [110/20] via 10.1.34.3, 00:04:55, Serial0/0.43
O N2 3.0.0.0/8 [110/20] via 10.1.34.3, 00:04:55, Serial0/0.43
O IA    5.5.5.5 [110/129] via 10.1.34.3, 00:05:01, Serial0/0.43
O IA    10.1.12.0 [110/192] via 10.1.34.3, 00:05:05, Serial0/0.43
O IA    10.1.23.0 [110/128] via 10.1.34.3, 00:05:05, Serial0/0.43
O IA    10.1.35.0 [110/128] via 10.1.34.3, 00:05:05, Serial0/0.43
O*N2 0.0.0.0/0 [110/1] via 10.1.34.3, 00:04:55, Serial0/0.43
```

Note the default route is injected as an external route.

On R5

R5#Show ip route ospf | Inc O

```
O IA    1.1.1.1 [110/193] via 10.1.35.3, 00:07:14, Serial0/0.53
O IA    2.2.2.2 [110/129] via 10.1.35.3, 00:07:14, Serial0/0.53
O N2 33.0.0.0/8 [110/20] via 10.1.35.3, 00:07:14, Serial0/0.53
O N2 3.0.0.0/8 [110/20] via 10.1.35.3, 00:07:14, Serial0/0.53
O IA    4.4.4.4 [110/129] via 10.1.35.3, 00:07:14, Serial0/0.53
O IA    10.1.12.0 [110/192] via 10.1.35.3, 00:07:14, Serial0/0.53
O IA    10.1.23.0 [110/128] via 10.1.35.3, 00:07:14, Serial0/0.53
O IA    10.1.34.0 [110/128] via 10.1.35.3, 00:07:14, Serial0/0.53
O*N2 0.0.0.0/0 [110/1] via 10.1.35.3, 00:07:14, Serial0/0.53
```

Task 4

Configure area 1 such that it receives the default route injected by the ABR as an internal OSPF route.

On R3

R3(config)#router ospf 1
R3(config-router)#area 1 nssa no-summary

To Verify the configuration:

On R4

R4#Sh ip rou ospf | inc O

O N2 33.0.0.0/8 [110/20] via 10.1.34.3, 00:12:18, Serial0/0.43
O N2 3.0.0.0/8 [110/20] via 10.1.34.3, 00:12:18, Serial0/0.43
O*IA 0.0.0.0/0 [110/65] via 10.1.34.3, 00:00:09, Serial0/0.43

Note the default route injected by the ABR of this area is an internal OSPF route.

Task 5

Configure R3 such that ONLY Area 2 receives the redistributed routes (3.0.0.0 /8 and 33.0.0.0 /8); you should NOT use any global configuration command or route-map as part of the solution to accomplish this task.

On R3

R3(config)#router ospf 1
R3(config-router)#area 1 nssa no-redistribution

To verify the configuration:

On R4

R4#Sh ip route ospf | Inc O

O*IA 0.0.0.0/0 [110/65] via 10.1.34.3, 00:01:45, Serial0/0.43

Note the no-redistribution is configured on the ABR which happens to be an ASBR as well; this command stops redistribution of the external routes into the area specified.

Task 6

Erase the startup configuration and reload the routers before proceeding to the next lab.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

BGP

Lab 1 – Establishing Neighbor Adjacency

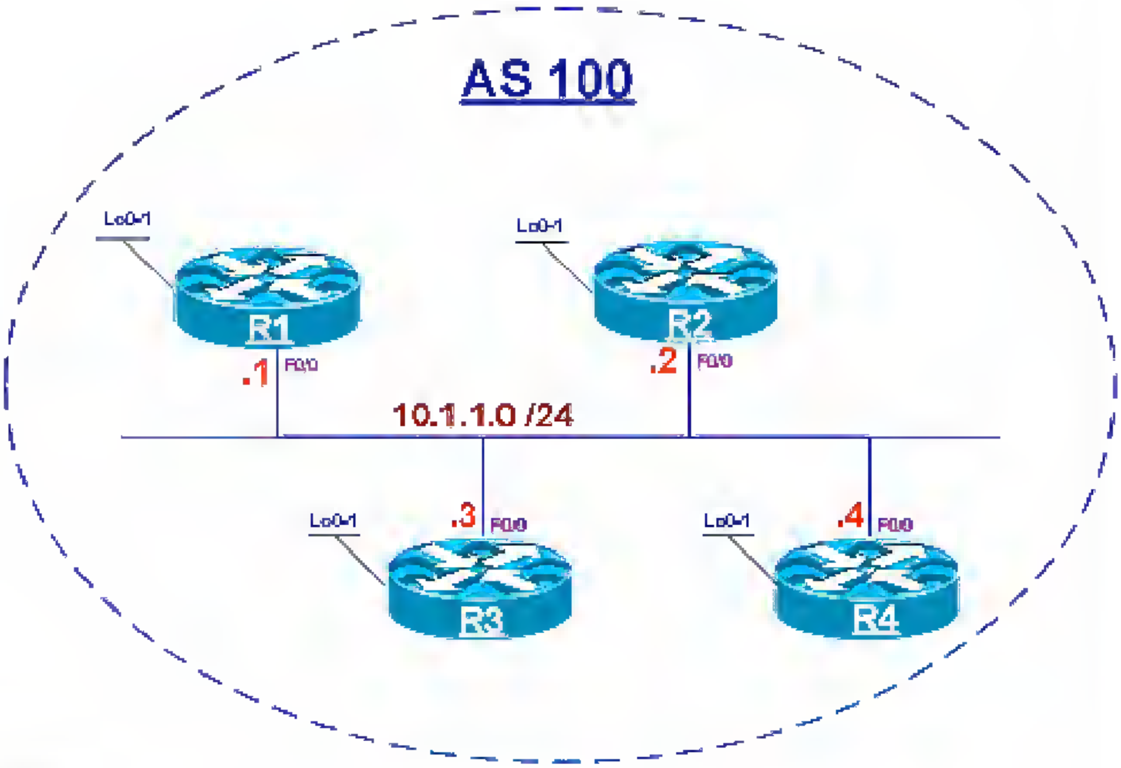

Lab Setup:

- The F0/0 interface of these four routers should be configured in VLAN 100.
- Configure the routers according to the following IP addressing chart:

IP Addressing:

Router	Interface / IP Address	AS 100
R1	F0/0 – 10.1.1.1 /24 Lo0 – 1.1.1.1 /8 Lo1 – 192.168.1.1 /24	100
R2	F0/0 – 10.1.1.2 /24 Lo0 – 2.2.2.2 /8 Lo1 – 192.168.2.2 /24	100

R3	F0/0 – 10.1.1.3 /24 Lo0 – 3.3.3.3 /8 Lo1 – 192.168.3.3 /24	100
R4	F0/0 – 10.1.1.4 /24 Lo0 – 4.4.4.4 /8 Lo1 – 192.168.4.4 /24	100

Task 1

Configure these routers in AS 100, these routers should create an IBGP peer sessions between them, ensure that these routers advertise their Loopback 0 interface in this AS.

On R1

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote-as 100
R1(config-router)#neighbor 10.1.1.3 remote-as 100
R1(config-router)#neighbor 10.1.1.4 remote-as 100
R1(config-router)#no syn
R1(config-router)#network 1.0.0.0
```

On R2

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.1.1 remote-as 100
R2(config-router)#neighbor 10.1.1.3 remote-as 100
R2(config-router)#neighbor 10.1.1.4 remote-as 100
R2(config-router)#no syn
R2(config-router)#network 2.0.0.0
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.1.1 remote-as 100
R3(config-router)#neighbor 10.1.1.2 remote-as 100
R3(config-router)#neighbor 10.1.1.4 remote-as 100
R3(config-router)#no syn
R3(config-router)#network 3.0.0.0
```

On R4

```
R1(config)#router bgp 100
```

```
R4(config-router)#neighbor 10.1.1.1 remote-as 100
R4(config-router)#neighbor 10.1.1.2 remote-as 100
R4(config-router)#neighbor 10.1.1.3 remote-as 100

R4(config-router)#no syn
R4(config-router)#network 4.0.0.0
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          10.1.1.2                 0    100      0 i
*>i3.0.0.0          10.1.1.3                 0    100      0 i
*>i4.0.0.0          10.1.1.4                 0    100      0 i
```

*	Valid Table Entry.
>	The best entry for the Prefix.
s	The entry is suppressed.
i	The entry was learned via an IBGP, this is the "i" to the left of the network column. The letter "i" under the path column, specifies the origin of the route.
Network	Prefix entry for the network, if the mask is omitted, the default mask is assumed.
Next Hop	The next hop's IP address to get to the specified network address, if it is 0.0.0.0 it is a prefix that is advertised by the local router.
Metric	This is the Inter-as metric, or the MED attribute which is 0 by default.
LocPrf	This is the local preference attribute, used in the route selection process carried within the local AS ONLY. With the local-pref attribute the higher value has more preference. The prefixes that are received from a peer AS are tagged with a local-pref value of 100; this value can be changed to influence the best path selection process. The changed value is only advertised to IBGP peers. When the local router advertises a prefix, no local-pref value is seen in the output of the "Show ip bgp" command. The default value of 100 can be changed by the **"BGP default local-preference"** command.
Weight	The prefixes that are received via a neighbor (IBGP or EBGP) will have a weight of 0, but the prefixes that are originated by the local router will have a weight value of 32768. This attribute overrides any other attribute for performing best path determination.
Path	If the prefixes were originated or learned via an IBGP neighbor, the path column will have the letter "i" without any ASN. If the prefix was learned through another AS, then this column will have the AS number/s followed by the letter i, the ASNs indicate the ASes that a prefix has traversed. The maximum number of ASes that a prefix can traverse through is 255.

Task 2

Reconfigure the routers as follows:

R2, R3 and R4 should be configured in AS 200, 300 and 400 respectively. Configure a full mesh peer session between these routers.

On R1

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote-as 200
R1(config-router)#neighbor 10.1.1.3 remote-as 300
R1(config-router)#neighbor 10.1.1.4 remote-as 400

```
R1(config-router)#no auto
R1(config-router)#network 1.0.0.0
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.1.1 remote-as 100
R2(config-router)#neighbor 10.1.1.3 remote-as 300
R2(config-router)#neighbor 10.1.1.4 remote-as 400
R2(config-router)#no auto
R2(config-router)#network 2.0.0.0
```

On R3

```
R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.1.1 remote-as 100
R3(config-router)#neighbor 10.1.1.2 remote-as 200
R3(config-router)#neighbor 10.1.1.4 remote-as 400
R3(config-router)#no auto
R3(config-router)#network 3.0.0.0
```

On R4

```
R1(config)#router bgp 400
R4(config-router)#neighbor 10.1.1.1 remote-as 100
R4(config-router)#neighbor 10.1.1.2 remote-as 200
R4(config-router)#neighbor 10.1.1.3 remote-as 300
R4(config-router)#no auto
R4(config-router)#network 4.0.0.0
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 200 i
```

```
*               10.1.1.2                                0 300 200 i
*               10.1.1.2                                0 400 200 i
* 3.0.0.0       10.1.1.3                                0 200 300 i
*>              10.1.1.3          0                     0 300 i
*               10.1.1.3                                0 400 300 i
* 4.0.0.0       10.1.1.4                                0 200 400 i
*               10.1.1.4                                0 300 400 i
*>              10.1.1.4          0                     0 400 i
```

Note the local-preference attribute is not assigned on any of the prefixes; this is because the prefixes are advertised by an EBGP peer. The best selection in the above output is based on the shortest ASN.
The MED value (Metric column) is zero for some of the prefixes, and on others, it is NOT assigned, this is because when the prefix is advertised by the originating AS, the metric is set to "0", but when the same prefix is advertised by another AS, the MED value is removed.

Task 3

Reconfigure the routers in AS 100; use the following policy for their IBGP peer sessions:

- Authentication must be enabled between the peers using "cisco" as the password.
- The peer session must be established based on the Loopback 0's IP address.
- These routers should ONLY advertise their Loopback 1 in BGP.
- Provide NLRI to Loopback0 interface using RIPv2.
- The peer session between the routers should only be established if they are running BGP version 4.
- Use peer-groups to accomplish this task.

Cisco's implementation of BGP in IOS 12.0(5)T or earlier releases supports BGP versions 2, 3, and 4, with dynamic negotiation down to Version 2. But in IOS version 12.0(6)T or later, Cisco routers only support version 4 and they do not support dynamic negotiation down to Version 2. The reason you may see the "Neighbor version" command configured on some Cisco routers is because may be the router is connecting and establishing a peer session with a Non-Cisco router, or the administrator is not aware of this fact.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no au
```

R1(config-router)#net 192.168.1.0

R1(config-router)#neighbor TST peer-group
R1(config-router)#neighbor TST remote-as 100
R1(config-router)#neighbor TST update-source lo0
R1(config-router)#neighbor TST version 4
R1(config-router)#neighbor TST password cisco

R1(config-router)#neighbor 2.2.2.2 peer-group TST
R1(config-router)#neighbor 3.3.3.3 peer-group TST
R1(config-router)#neighbor 4.4.4.4 peer-group TST

R1(config-router)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#netw 10.0.0.0
R1(config-router)#netw 1.0.0.0

On R2

R2(config)#router bgp 100
R2(config-router)#no au
R2(config-router)#netw 192.168.2.0

R2(config-router)#neighbor TST peer-group
R2(config-router)#neighbor TST remote-as 100
R2(config-router)#neighbor TST update-source lo0
R2(config-router)#neighbor TST version 4
R2(config-router)#neighbor TST password cisco

R2(config-router)#neighbor 1.1.1.1 peer-group TST
R2(config-router)#neighbor 3.3.3.3 peer-group TST
R2(config-router)#neighbor 4.4.4.4 peer-group TST

R2(config-router)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
R2(config-router)#netw 2.0.0.0

On R3

R3(config)#router bgp 100
R3(config-router)#no au

```
R3(config-router)# network 192.168.3.0

R3(config-router)#neighbor TST peer-group
R3(config-router)#neighbor TST remote-as 100
R3(config-router)#neighbor TST update-source lo0
R3(config-router)#neighbor TST version 4
R3(config-router)#neighbor TST password cisco

R3(config-router)#neighbor 1.1.1.1 peer-group TST
R3(config-router)#neighbor 2.2.2.2 peer-group TST
R3(config-router)#neighbor 4.4.4.4 peer-group TST

R3(config-router)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
R3(config-router)#netw 3.0.0.0
```

On R4

```
R4(config)#router bgp 100
R4(config-router)#no au
R4(config-router)# network 192.168.4.0

R4(config-router)#neighbor TST peer-group
R4(config-router)#neighbor TST remote-as 100
R4(config-router)#neighbor TST update-source lo0
R4(config-router)#neighbor TST version 4
R4(config-router)#neighbor TST password cisco

R4(config-router)#neighbor 1.1.1.1 peer-group TST
R4(config-router)#neighbor 2.2.2.2 peer-group TST
R4(config-router)#neighbor 3.3.3.3 peer-group TST

R4(config-router)#router rip
R4(config-router)#no au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
R4(config-router)#netw 4.0.0.0
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop     Metric  LocPrf  Weight  Path
*> 192.168.1.0      0.0.0.0           0           32768  i
*>i192.168.2.0      2.2.2.2           0     100       0  i
*>i192.168.3.0      3.3.3.3           0     100       0  i
*>i192.168.4.0      4.4.4.4           0     100       0  i
```

To verify the configuration:

On R1

R1#Show ip bgp peer-group

```
BGP peer-group is TST, remote AS 100
  BGP version 4
  Default minimum time between advertisement runs is 0 seconds

 For address family: IPv4 Unicast
  BGP neighbor is TST, peer-group internal, members:
  2.2.2.2 3.3.3.3 4.4.4.4
  Index 0, Offset 0, Mask 0x0
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
```

Note the output of the "Show ip bgp peer-group" reveals the ip address of the members of the peer-group.

Some of the benefits of peer-groups:

- **Peer-groups provide optimization of BGP convergence, Let's say a BGP speaker has 10 IBGP peers that exchange full BGP routing (200,000 prefixes), without the creation of a peer-group, the local router has to go through 2 million prefixes, whereas, if the same router was configured with a peer-group, the router would only go through 200,000 prefixes.**

- **It provides a mechanism for peers that have an identical outbound policy.**

> **Another benefit of peer-groups is that it can reduce the administrative overhead by cutting down redundant configuration on the routers.**

Task 4

Remove the BGP configuration from the routers and reconfigure the routers in AS 100 using peer-session templates; you should configure the following two templates to accomplish this task:

- Common Template: This template should contain the "**Neighbor version 4**" and "**Neighbor password**" command, this template should be applied to all neighbors.

- IBGP Template: This template should contain the "**Neighbor Update-source**" and "**Neighbor remote-as**" commands. This template should be applied to all IBGP neighbors.

You should advertise Loopback1 interface in BGP and Loopback0 should be used as the IP address for establishing the peer sessions. DO NOT remove RIPv2's configuration.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no au
R1(config-router)#network 192.168.1.0

R1(config-router)#template peer-session Common
R1(config-router-stmp)#password cisco
R1(config-router-stmp)#version 4
R1(config-router-stmp)#exit-peer-session

R1(config-router)#template peer-session IBGP
R1(config-router-stmp)#inherit peer-session Common
R1(config-router-stmp)#update-source lo0
R1(config-router-stmp)#remote-as 100
R1(config-router-stmp)#exit-peer-session

R1(config-router)#neighbor 2.2.2.2 inherit peer-session IBGP
R1(config-router)#neighbor 3.3.3.3 inherit peer-session IBGP
R1(config-router)#neighbor 4.4.4.4 inherit peer-session IBGP
```

On R2

```
R2(config)#router bgp 100
R2(config-router)#no au
R2(config-router)#network 192.168.2.0

R2(config-router)#template peer-session Common
R2(config-router-stmp)#password cisco
R2(config-router-stmp)#version 4
R2(config-router-stmp)#exit-peer-session

R2(config-router)#template peer-session IBGP
R2(config-router-stmp)#inherit peer-session Common
R2(config-router-stmp)#update-source lo0
R2(config-router-stmp)#remote-as 100
R2(config-router-stmp)#exit-peer-session

R2(config-router)#neighbor 1.1.1.1 inherit peer-session IBGP
R2(config-router)#neighbor 3.3.3.3 inherit peer-session IBGP
R2(config-router)#neighbor 4.4.4.4 inherit peer-session IBGP
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#no au
R3(config-router)#network 192.168.3.0

R3(config-router)#template peer-session Common
R3(config-router-stmp)#password cisco
R3(config-router-stmp)#version 4
R3(config-router-stmp)#exit-peer-session

R3(config-router)#template peer-session IBGP
R3(config-router-stmp)#inherit peer-session Common
R3(config-router-stmp)#update-source lo0
R3(config-router-stmp)#remote-as 100
R3(config-router-stmp)#exit-peer-session

R3(config-router)#neighbor 1.1.1.1 inherit peer-session IBGP
R3(config-router)#neighbor 2.2.2.2 inherit peer-session IBGP
R3(config-router)#neighbor 4.4.4.4 inherit peer-session IBGP
```

On R4

```
R4(config)#router bgp 100
R4(config-router)#no au
R4(config-router)#network 192.168.4.0

R4(config-router)#template peer-session Common
R4(config-router-stmp)#password cisco
R4(config-router-stmp)#version 4
R4(config-router-stmp)#exit-peer-session

R4(config-router)#template peer-session IBGP
R4(config-router-stmp)#inherit peer-session Common
R4(config-router-stmp)#update-source lo0
R4(config-router-stmp)#remote-as 100
R4(config-router-stmp)#exit-peer-session

R4(config-router)#neighbor 1.1.1.1 inherit peer-session IBGP
R4(config-router)#neighbor 2.2.2.2 inherit peer-session IBGP
R4(config-router)#neighbor 3.3.3.3 inherit peer-session IBGP
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 5, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop         Metric  LocPrf Weight  Path
*> 192.168.1.0      0.0.0.0               0          32768  i
*>i192.168.2.0      2.2.2.2               0     100      0  i
*>i192.168.3.0      3.3.3.3               0     100      0  i
*>i192.168.4.0      4.4.4.4               0     100      0  i
```

- **Peer-session template can be used to apply session specific configuration commands to a group of neighbors that share a common session configuration.**

- **Peer-session templates can be reused and they support inheritance of another peer-session template/s, this means that nested peer-sessions can also be used.**

> ➢ Peer-session templates support session specific commands ONLY.

Task 5

Reconfigure the routers based on the following IP address space and diagram.

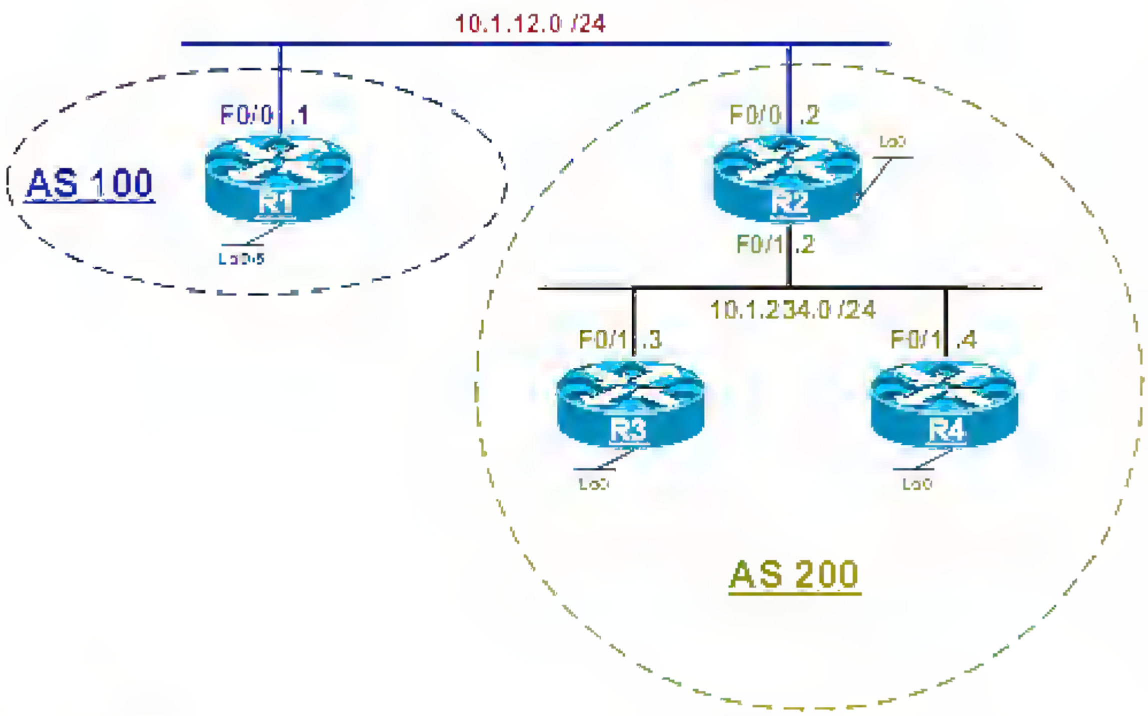

Lab Setup:

- ➢ Configure F0/0 interface of R1 and R2 are in VLAN 12 and the F0/1 interface of R2, R3 and R4 in VLAN 234.
- ➢ Configure IP addressing on the routers using the IP addressing chart on the next page.

IP Addressing chart:

Router	AS number	Interface / IP address
R1	AS 100	Lo0 = 1.1.0.1 /24 Lo1 = 1.1.1.1 /24 Lo2 = 1.1.2.1 /24 Lo3 = 1.1.3.1 /24 Lo4 = 100.1.1.1 /24 Lo5 = 100.2.2.1 /24 F0/0 = 10.1.12.1 /24
R2	AS 200	Lo0 = 2.2.2.2 /8 F0/0 = 10.1.12.2 /24 F0/1 = 10.1.234.2 /24
R3	AS 200	Lo0 = 3.3.3.3 /8 F0/1 = 10.1.234.3 /24
R4	AS 200	Lo0 = 4.4.4.4 /8 F0/1 = 10.1.234.4 /24

Task 6

- R1 in AS 100 should establish an EBGP peer session with R2 in AS 200. R1 should advertise all of it's loopback interfaces in AS 100.

- R2, R3 and R4 should be configured in AS 200; these routers should establish IBGP peer sessions between them and advertise their loopback 0 interface in AS 200.
- Configure the router-ids of the routers as follows:

  R1 = 10.1.1.1, R2 = 10.2.2.2, R3 = 10.3.3.3 and R4 = 10.4.4.4

On R1

```
R1(config)#router bgp 100
R1(config-router)#bgp router-id 10.1.1.1
R1(config-router)#network 1.1.0.0 mask 255.255.255.0
R1(config-router)#network 1.1.1.0 mask 255.255.255.0
R1(config-router)#network 1.1.2.0 mask 255.255.255.0
R1(config-router)#network 1.1.3.0 mask 255.255.255.0
R1(config-router)#network 100.1.1.0 mask 255.255.255.0
R1(config-router)#network 100.2.2.0 mask 255.255.255.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200

R1(config-router)#no auto-summary
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no synchronization
R2(config-router)#bgp router-id 10.2.2.2
R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.234.3 remote-as 200
R2(config-router)#neighbor 10.1.234.4 remote-as 200
R2(config-router)#no auto-summary
```

On R3

```
R3(config)#router bgp 200
R3(config-router)#no synchronization
R3(config-router)#bgp router-id 10.3.3.3
R3(config-router)#nework 3.0.0.0
R3(config-router)#neighbor 10.1.234.2 remote-as 200
R3(config-router)#neighbor 10.1.234.4 remote-as 200
R3(config-router)#no auto-summary
```

On R4

```
R4(config)#router bgp 200
R4(config-router)#no synchronization
R4(config-router)#bgp router-id 10.4.4.4
R4(config-router)#nework 4.0.0.0
R4(config-router)#neighbor 10.1.234.2 remote-as 200
R4(config-router)#neighbor 10.1.234.3 remote-as 200
R4(config-router)#no auto-summary
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 16, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       0.0.0.0                  0         32768 i
```

```
*> 1.1.1.0/24       0.0.0.0            0         32768 i
*> 1.1.2.0/24       0.0.0.0            0         32768 i
*> 1.1.3.0/24       0.0.0.0            0         32768 i
*> 2.0.0.0          10.1.12.2          0             0 200 i
*> 3.0.0.0          10.1.12.2                        0 200 i
*> 4.0.0.0          10.1.12.2                        0 200 i
*> 100.1.1.0/24     0.0.0.0            0         32768 i
*> 100.2.2.0/24     0.0.0.0            0         32768 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 10, local router ID is 10.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop       Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1          0             0 100 i
*> 1.1.1.0/24       10.1.12.1          0             0 100 i
*> 1.1.2.0/24       10.1.12.1          0             0 100 i
*> 1.1.3.0/24       10.1.12.1          0             0 100 i
*> 2.0.0.0          0.0.0.0            0         32768 i
*>i3.0.0.0          10.1.234.3         0    100      0 i
*>i4.0.0.0          10.1.234.4         0    100      0 i
*>100.1.1.0/24      10.1.12.1          0             0 100 i
*> 100.2.2.0/24     10.1.12.1          0             0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 18, local router ID is 10.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop       Metric LocPrf Weight Path
* i1.1.0.0/24       10.1.12.1          0    100      0 100 i
* i1.1.1.0/24       10.1.12.1          0    100      0 100 i
* i1.1.2.0/24       10.1.12.1          0    100      0 100 i
* i1.1.3.0/24       10.1.12.1          0    100      0 100 i
*>i2.0.0.0          10.1.234.2         0    100      0 i
```

```
*> 3.0.0.0          0.0.0.0              0          32768  i
*>i4.0.0.0          10.1.234.4           0    100       0 i
* i100.1.1.0/24     10.1.12.1            0    100       0 100 i
* i100.2.2.0/24     10.1.12.1            0    100       0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 20, local router ID is 10.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight  Path
* i1.1.0.0/24       10.1.12.1            0     100       0   100 i
* i1.1.1.0/24       10.1.12.1            0     100       0   100 i
* i1.1.2.0/24       10.1.12.1            0     100       0   100 i
* i1.1.3.0/24       10.1.12.1            0     100       0   100 i
*>i2.0.0.0          10.1.234.2           0     100       0    i
*>i3.0.0.0          10.1.234.3           0     100       0    i
*> 4.0.0.0          0.0.0.0              0           32768    i
* i100.1.1.0/24     10.1.12.1            0     100       0   100 i
* i100.2.2.0/24     10.1.12.1            0     100       0   100 i
```

Note R3 and R4 do not have NLRI to the next-hop IP address of 10.1.12.1, therefore, they won't have reachability to these addresses.

Task 7

Configure R2 to change the next hop IP address for all the networks advertised by R1 to the IP address of it's F0/1 interface. You should use a template so the future policies can be installed once in that template and have it effect R3 and R4. DO NOT use peer-groups to accomplish this task.

On R2

```
R2(config)#router bgp 200
R2(config-router)# template peer-policy TST
R2(config-router-ptmp)# next-hop-self
R2(config-router-ptmp)#exit-peer-policy
```

R2(config-router)#neighbor 10.1.234.3 inherit peer-policy TST
R2(config-router)#neighbor 10.1.234.4 inherit peer-policy TST

To verify the configuration:

On R3

R3#Sh ip bgp

```
BGP table version is 18, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight  Path
*> i1.1.0.0/24      10.1.234.2           0    100      0  100 i
*> i1.1.1.0/24      10.1.234.2           0    100      0  100 i
*>i1.1.2.0/24       10.1.234.2           0    100      0  100 i
*> i1.1.3.0/24      10.1.234.2           0    100      0  100 i
*>i2.0.0.0          10.1.234.2           0    100      0  i
*> 3.0.0.0          0.0.0.0              0         32768  i
*>i4.0.0.0          10.1.234.4           0    100      0  i
*> i100.1.1.0/24    10.1.234.2           0    100      0  100 i
*>i100.2.2.0/24     10.1.234.2           0    100      0  100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 20, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight  Path
*> i1.1.0.0/24      10.1.234.2           0    100      0  100 i
*> i1.1.1.0/24      10.1.234.2           0    100      0  100 i
*> i1.1.2.0/24      10.1.234.2           0    100      0  100 i
*> i1.1.3.0/24      10.1.234.2           0    100      0  100 i
*>i2.0.0.0          10.1.234.2           0    100      0  i
*>i3.0.0.0          10.1.234.3           0    100      0  i
*> 4.0.0.0          0.0.0.0              0         32768  i
*>i100.1.1.0/24     10.1.234.2           0    100      0  100 i
*> i100.2.2.0/24    10.1.234.2           0    100      0  100 i
```

Peer-policy templates are used to build a template of policy information that can be inherited by a given neighbor. The peer-policy template can not be inherited by a peer-session template or a peer-group.

Task 8

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 2
Route Reflectors

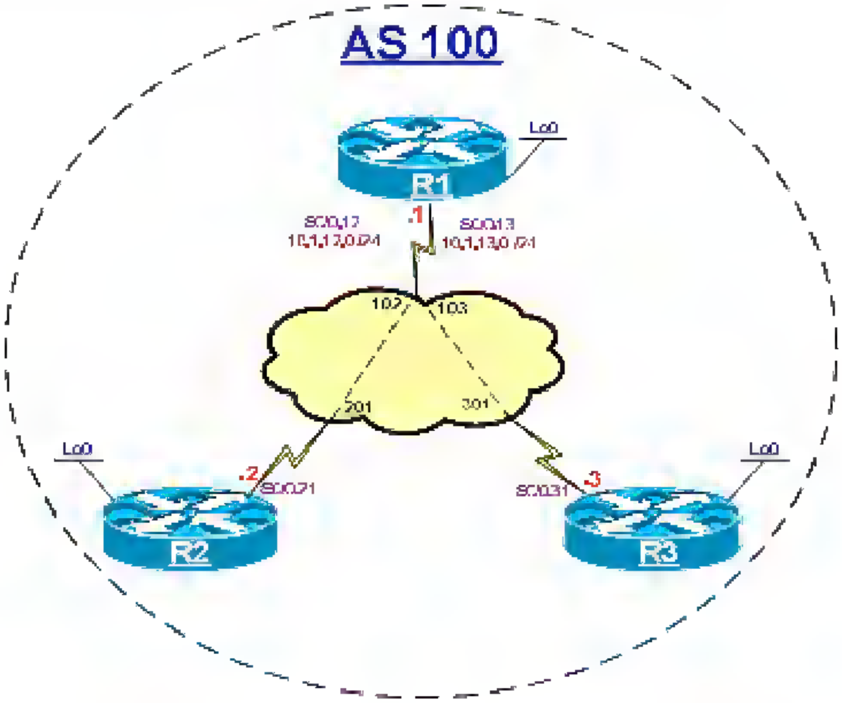

Lab Setup:

- Configure R1 to have two point-to-point frame-relay connections, one connecting R1 to R2, and the other connecting R1 to R3.
- R2 and R3 should each be configured with a frame-relay point-to-point connection to R1
- Use the following IP address chart for IP address assignment.

IP Addressing:

Router	Interface	IP address
R1	Lo0 The frame-relay connection to R2 The frame-relay connection to R3	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24
R2	Lo0 The frame-relay connection to R1	2.2.2.2 /8 10.1.12.2 /24
R3	Lo0 The frame-relay connection to R1	3.3.3.3 /8 10.1.13.3 /24

Task 1

Configure BGP AS 100 on all routers and ensure that the routers can successfully establish an IBGP peer session with each other. These routers should only advertise their Loopback0 interface in BGP. To provide NLRI, the links between the routers should be advertised in RIPv2.

On All Routers

```
(config-router)#router rip
(config-router)#no au
(config-router)#ver 2
(config-router)#netw 10.0.0.0
```

On R1

```
R1(config)#router bgp 100
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 100
R1(config-router)#neighbor 10.1.13.3 remote-as 100
R1(config-router)#no syn
```

On R2

```
R2(config)#router bgp 100
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.13.3 remote-as 100
R2(config-router)#no syn
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.13.1 remote-as 100
R3(config-router)#neighbor 10.1.12.2 remote-as 100
R3(config-router)#no syn
```

To verify the configuration:

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 10.1.13.1, 00:20:04
B    2.0.0.0/8 [200/0] via 10.1.12.2, 00:20:09
C    3.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
R       10.1.12.0 [120/1] via 10.1.13.1, 00:00:25, Serial0/0.31
```

On R2

R2#Show ip bgp

```
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
```

Task 2

You received an e-mail from the management stating that within the next 12 months 20 additional routers will be added to this AS. In order to minimize the number of peer sessions within this AS, you decided to implement route reflectors. Configure R1 as a route reflector for this AS.

On R1

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-reflector-client
R1(config-router)#neighbor 10.1.13.3 route-reflector-client
```

On R2

```
R2(config)#router bgp 100
R2(config-router)#no neighbor 10.1.13.3 remote-as 100
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#NO neighbor 10.1.12.2 remote-as 100
```

In order for all IBGP speakers in an AS to exchange routes with one another, the IBGP speakers must be fully meshed (Every router must establish a peer session to every other router). Route-reflectors can be configured to reduce the number of peer sessions that must be established between the routers within a given AS. If a route-reflector is used, all IBGP speakers need not be fully meshed. In this model, the router that is configured to be the route-reflector must have a peer session established to every client, the clients must establish a peer session with the route reflector. The route reflector will reflect routes learned from one client to the other client/s.

To verify the configuration:

On R2

```
R2#Show ip bgp

BGP table version is 10, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight  Path
*>i1.0.0.0          10.1.12.1                0    100      0  i
*> 2.0.0.0          0.0.0.0                  0         32768  i
*>i3.0.0.0          10.1.13.3                0    100      0  i
```

R2#Ping 3.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/113/116 ms
```

On R3

R3#Show ip bgp

```
BGP table version is 10, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric  LocPrf Weight  Path
*>i1.0.0.0          10.1.13.1                0     100      0  i
*>i2.0.0.0          10.1.12.2                0     100      0  i
*> 3.0.0.0          0.0.0.0                  0          32768  i
```

R3#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/112/116 ms
```

Task 3

After implementing the route reflector, you realized that if the route reflector is down, the entire network is dysfunctional; therefore, you decided to add R4 as the second route reflector for redundancy. Ensure that these routers can reach the advertised networks and the redundancy is operational.

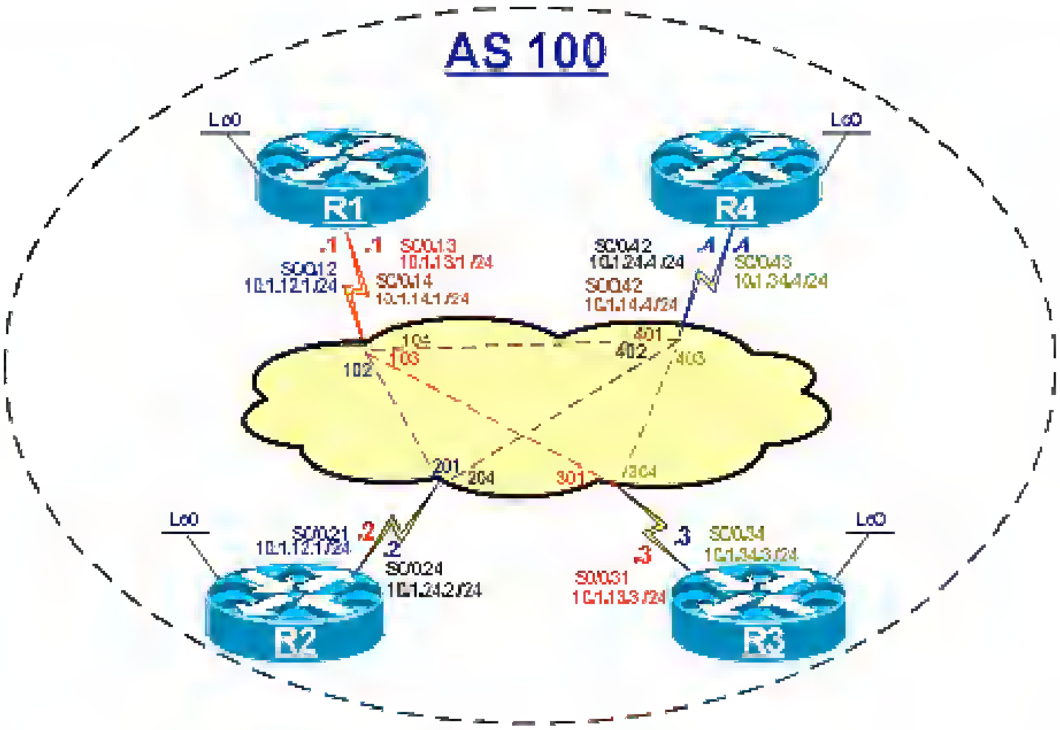

DO NOT erase the existing configuration, the following configuration is added to the existing configuration.

Lab Setup:

Add the following configuration to the existing configuration:

- Configure R1 with an additional point-to-point frame-relay connection to R4, using the IP addressing and the DLCI information provided below. R1 should establish a BGP peer session with R4 over this frame-relay connection.

- R2 and R3 should each be configured with an additional point-to-point frame-relay connection to R4. Use the IP addressing and the DLCI information provided below for these connections. R2 and R3 should each establish a BGP peer session with R4 over this connection.
- R4 should be configured with three point-to-point frame-relay connections, one to R1, the second one to R2 and the third one to R3. Use the following IP addressing and DLCI information for these connections.

IP Addressing:

Router	Interface	IP address	DLCI
R1	S0/0.14	10.1.14.1 /24	104
R2	S0/0.24	10.1.24.2 /24	204
R3	S0/0.34	10.1.34.3 /24	304
R4	Lo0 S0/0.42 S0/0.43 S0/0.41	4.4.4.4 /8 10.1.24.4 /24 10.1.34.4 /24 10.1.14.4 /24	 402 403 401

On R4

```
R4(config)#router bgp 100
R4(config)#Network 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.24.2 remote-as 100
R4(config-router)#neighbor 10.1.34.3 remote-as 100
R4(config-router)#neighbor 10.1.24.2 route-reflector client
R4(config-router)#neighbor 10.1.34.3 route-reflector client
```

R4 is the secondary route-reflector. R4 should be configured as follows:

- R4 should have a peer session with R1 – the route-reflectors should have full mesh peer sessions between them.
- R4 must have a peer session with R2 and R3.
- R4 must configure R2 and R3 as route-reflector clients.

On R2

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.24.4 remote-as 100
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.34.4 remote-as 100
```

On R1

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.14.4 remote-as 100
```

Having a single RR can introduce a single point of failure, its best to have multiple RRs incase the RR fails, this redundancy is critical when there are many RR clients.

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
        r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
* i2.0.0.0          10.1.24.2                0    100      0 i
*>i                 10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
* i                 10.1.34.3                0    100      0 i
*>i4.0.0.0          10.1.14.4                0    100      0 i
```

To test the configuration:

On R1

```
R1(config)#int s0/0
R1(config-if)#shut
```

On R2

R2#Show ip bgp

```
BGP table version is 8, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
        r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*>i4.0.0.0          10.1.24.4                0    100      0 i
```

```
R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/114/116 ms
```

Task 4

Erase the startup configuration and reload the routers. Reconfigure the routers based on the following IP addressing and topology:

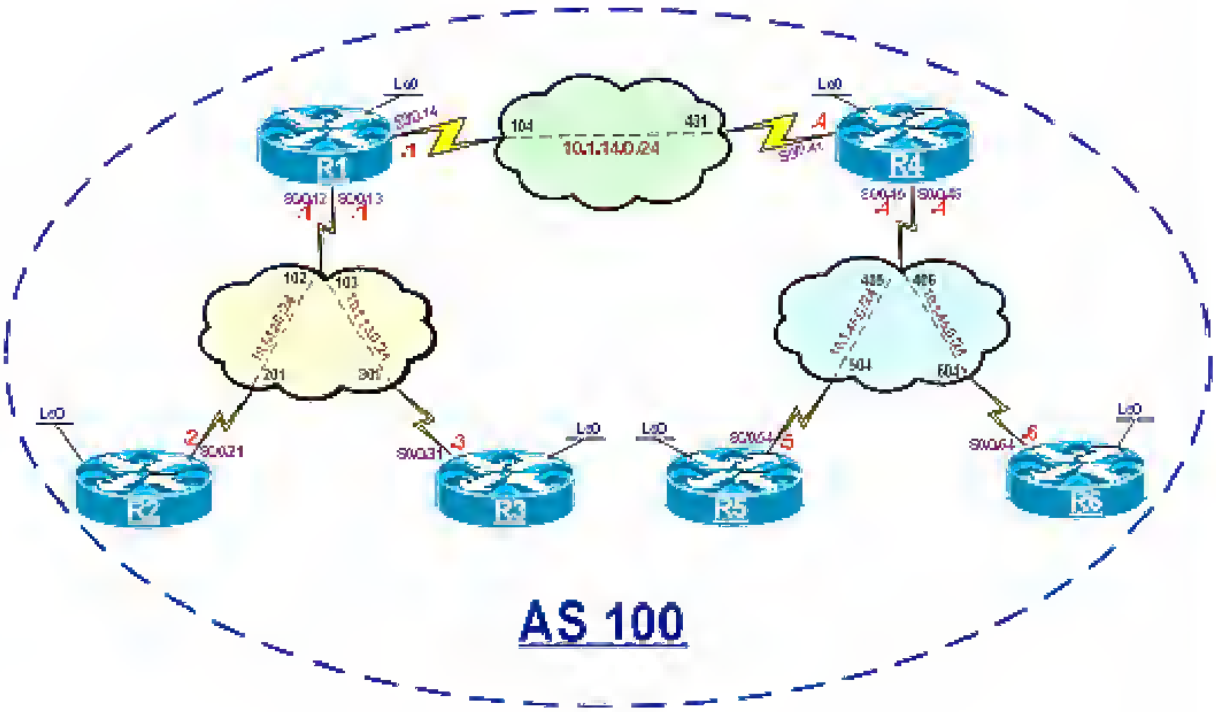

IP Addressing:

Router	Interface	IP address
R1	Lo0 The frame-relay connection to R4 The frame-relay connection to R2 The frame-relay connection to R3	1.1.1.1 /8 10.1.14.1 /24 10.1.12.1 /24 10.1.13.1 /24
R2	Lo0 The frame-relay connection to R1	2.2.2.2 /8 10.1.12.2 /24
R3	Lo0 The frame-relay connection to R1	3.3.3.3 /8 10.1.13.3 /24
R4	Lo0 The frame-relay connection to R1 The frame-relay connection to R5 The frame-relay connection to R6	4.4.4.4 /8 10.1.14.4 /24 10.1.45.4 /24 10.1.46.4 /24
R5	Lo0 The frame-relay connection to R1	5.5.5.5 /8 10.1.45.5 /24
R6	Lo0 The frame-relay connection to R1	6.6.6.6 /8 10.1.46.6 /24

Lab Setup:

- Configure R1 with three point-to-point frame-relay connections; these point-to-point connections should connect R1 to R2, R3 and R4.
- Configure R2 and R3 with a single point-to-point connection to R1.
- R4 should be configured with three point-to-point frame-relay connections; these point-to-point connections should connect R4 to R5, R6 and R1.
- Configure R5 and R6 with a single point-to-point connection to R4.
- R1 should be configured as the route reflector for routers R2 and R3, whereas R4 should be configured to be the route-reflector for routers R5 and R6.
- R1 and R4 should be configured to have an IBGP peer session between them; these two routers should be configured in BGP AS 100.
- NLRI for the links should be provided through RIPv2.

On All Routers

```
(config)#router rip
(config-router)#no au
(config-router)#ver 2
(config-router)#netw 10.0.0.0
```

On R1

```
R1(config-router)#router bgp 100
R1(config-router)#no au
R1(config-router)#no syn

R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 100
R1(config-router)#neighbor 10.1.13.3 remote-as 100
R1(config-router)#neighbor 10.1.14.4 remote-as 100
R1(config-router)#neighbor 10.1.12.2 route-reflector-client
R1(config-router)#neighbor 10.1.13.3 route-reflector-client
```

On R2

```
R2(config)#router bgp 100
R2(config-router)#no au
R2(config-router)#no syn

R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#no au
R3(config-router)#no syn
R3(config-router)#network 3.0.0.0
R3(config-router)#neighbor 10.1.13.1 remote-as 100
```

To verify the configuration:

On R3

R3#Show ip bgp

```
BGP table version is 4, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.13.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

On R4

```
R4(config)#router bgp 100
R4(config-router)#no au
R4(config-router)#no syn
R4(config-router)#network 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.45.5 remote-as 100
R4(config-router)#neighbor 10.1.46.6 remote-as 100
R4(config-router)#neighbor 10.1.45.5 route-reflector-client
R4(config-router)#neighbor 10.1.46.6 route-reflector-client
```

On R5

```
R5(config)#Router bgp 100
R5(config-router)#No au
R5(config-router)#No syn
R5(config-router)#Network 5.0.0.0
R5(config-router)#Neighbor 10.1.45.4 remote-as 100
```

On R6

```
R6(config)#Router bgp 100
R6(config-router)#No au
R6(config-router)#No syn
R6(config-router)#Network 6.0.0.0
R6(config-router)#Neighbor 10.1.46.4 remote-as 100
```

To verify the configuration:

On R6

R6#Show ip bgp

BGP table version is 7, local router ID is 6.6.6.6

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
*>i4.0.0.0          10.1.46.4                0    100      0 i
*>i5.0.0.0          10.1.45.5                0    100      0 i
*> 6.0.0.0          0.0.0.0                  0         32768 i
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 10.1.13.1, 00:08:44
B    2.0.0.0/8 [200/0] via 10.1.12.2, 00:08:44
C    3.0.0.0/8 is directly connected, Loopback0
B    4.0.0.0/8 [200/0] via 10.1.14.4, 00:04:14
B    5.0.0.0/8 [200/0] via 10.1.45.5, 00:02:57
B    6.0.0.0/8 [200/0] via 10.1.46.6, 00:02:03
     10.0.0.0/24 is subnetted, 5 subnets
R       10.1.14.0 [120/1] via 10.1.13.1, 00:00:07, Serial0/0.31
C       10.1.13.0 is directly connected, Serial0/0.31
R       10.1.12.0 [120/1] via 10.1.13.1, 00:00:07, Serial0/0.31
R       10.1.46.0 [120/2] via 10.1.13.1, 00:00:08, Serial0/0.31
R       10.1.45.0 [120/2] via 10.1.13.1, 00:00:08, Serial0/0.31
```

Note AS 100 has two route reflectors, each route-reflector has it's own clients, when a given RR receives an update from one of it's clients, it advertises that prefix to the other RR/s, the other RR/s in turn advertise that prefix to their clients.
There are some additional optional non-transitive attributes that can be used when RRs are configured and they are: originator-id, cluster-id and cluster-list.

- Originator-id: This attribute is created by the RR; this is the router-id of the router that originated the prefix. It's created to avoid routing loops, a RR will NOT advertise a route back to the originator of the prefix and if the originator of a prefix receives an update with its own router-id, it will ignore that prefix.

- Cluster and Cluster-id: A RR/s and its clients are collectively known as a cluster, each cluster must be uniquely identified, and the cluster-id is typically the router-id of the RR unless specifically configured.

- Cluster-list: This attribute is analogous to AS-path attribute, and it keeps track of the cluster-ids in the same way that the AS-path attribute keeps track of the AS numbers. When the RR advertises a prefix to a non-client, it appends the cluster-id to that prefix's cluster-list, if a RR receives an update and sees its own cluster-id in the cluster-list, it will ignore that update.

To see the attributes:

On R1

R1#Show ip bgp 6.0.0.0

```
BGP routing table entry for 6.0.0.0/8, version 21
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  Local
    10.1.46.6 (metric 1) from 10.1.14.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 6.6.6.6, Cluster list: 4.4.4.4
```

Note prefix 6.0.0.0 is the originator of the prefix and it came from 4.4.4.4 (The cluster-list).

On R3

R3#Show ip bgp 6.0.0.0

```
BGP routing table entry for 6.0.0.0/8, version 25

Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer

  Local
```

```
    10.1.46.6 (metric 2) from 10.1.13.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 6.6.6.6, Cluster list: 1.1.1.1, 4.4.4.4
```

Note this prefix has gone through cluster-ids of 4.4.4.4 first and then it traversed through cluster-id 1.1.1.1 before it was received by the local router.

Note the originator-id is the router-id of the router that originated that prefix, the output of the following "Show" command reveals the router-id of the router that originated the route.

```
R6#Sh ip bgp

BGP table version is 7, local router ID is 6.6.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.13.3                0    100      0 i
*>i4.0.0.0          10.1.46.4                0    100      0 i
*>i5.0.0.0          10.1.45.5                0    100      0 i
*> 6.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 3
Conditional Advertisement & BGP Backdoor

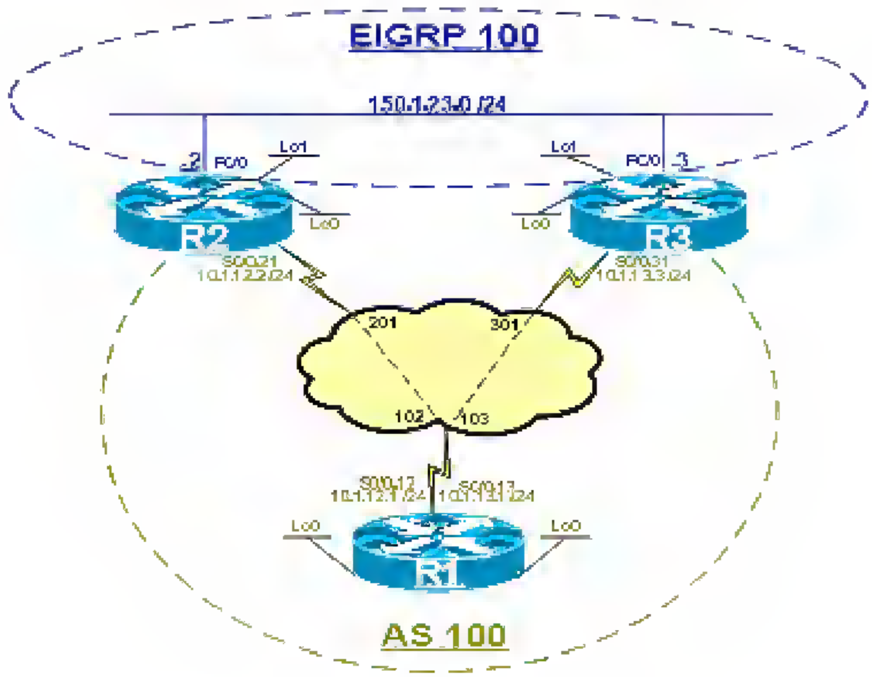

Lab Setup:

- Configure R1 to have two point-to-point frame-relay connections, one connecting R1 to R2, and the other connecting R1 to R3.
- R2 and R3 should be configured with a single frame-relay point-to-point connection to R1.
- Configure R2 and R3's F0/0 interface to be in VLAN 23.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12 S0/0.13	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24	100
R2	Lo0 Lo1 S0/0.21 F0/0	2.2.2.2 /8 150.1.2.2 /24 10.1.12.2 /24 150.1.23.2 /24	200
R3	Lo0 Lo1 S0/0.31 F0/0	3.3.3.3 /8 150.1.3.3 /24 10.1.13.3 /24 150.1.23.3 /24	300

Task 1

Configure EBGP peer sessions as follows:

- R1 to have an EBGP peer session to R2 and R3
- R2 and R3 to have an EBGP peer session to R1

On R1

```
R1(config)#router bgp 100
R1(config-router)#no auto
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.13.3 remote-as 300
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

On R3

```
R3(config-if)#router bgp 300
R3(config-router)#no au
R3(config-router)#neighbor 10.1.13.1 remote-as 100
```

To verify the configuration:

On R1

R1#Show ip bgp summary

```
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.12.2       4   200       3       3        1    0    0 00:00:56        0
10.1.13.3       4   300       4       4        1    0    0 00:00:20        0
```

Task 2

Configure R1, R2 and R3 to advertise their loopback0 interface in BGP.

On R1

```
R1(config-router)#router bgp 100
R1(config-router)#network 1.0.0.0
```

On R2

```
R2(config-router)#router bgp 200
R2(config-router)#network 2.0.0.0
```

On R3

```
R3(config)#router bgp 300
R3(config-router)#network 3.0.0.0
```

To verify the configuration:

On R3

R3#Show ip bgp

```
BGP table version is 4, local router ID is 150.1.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 3

Configure RIPv2 and Eigrp 100 on the routers as follows:

- Configure RIPv2 on all routers to advertise network 10.0.0.0, these routers should have their auto summarization disabled.
- R2 and R3 should also advertise their loopback1 and F0/0 interface in Eigrp AS#100.

On R2 and R3

```
(config)#router eigrp 100
(config-router)#no au
(config-router)#netw 150.1.0.0
```

On All Routers:

```
(config)#Router Rip
(config)#Ver 2
(config-router)#No au
(config-router)#Network 10.0.0.0
```

Task 4

If the link between R2 and R3 (The F0/0 interface) goes down, Loopback1 network of these two routers won't have connectivity even though there is a redundant link between these two routers, therefore, the administrator of R2 and R3 decided to advertise their Loopback 1 interface in BGP for redundancy, configure these routers to accommodate this decision.

On R2

R2(config-router)#router bgp 200
R2(config-router)#network 150.1.2.0 mask 255.255.255.0

On R3

R3(config)#router bgp 300
R3(config-router)#network 150.1.3.0 mask 255.255.255.0

To verify the configuration:

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [20/0] via 10.1.12.1, 00:18:54
C    2.0.0.0/8 is directly connected, Loopback0
B    3.0.0.0/8 [20/0] via 10.1.12.1, 00:18:54
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.13.0 [120/1] via 10.1.12.1, 00:00:01, Serial0/0.21
C       10.1.12.0 is directly connected, Serial0/0.21
     150.1.0.0/24 is subnetted, 3 subnets
C       150.1.23.0 is directly connected, FastEthernet0/0
B       150.1.3.0 [20/0] via 10.1.12.1, 00:13:21
C       150.1.2.0 is directly connected, Loopback1
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [20/0] via 10.1.13.1, 00:20:24
B    2.0.0.0/8 [20/0] via 10.1.13.1, 00:20:55
C    3.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
R       10.1.12.0 [120/1] via 10.1.13.1, 00:00:20, Serial0/0.31
     150.1.0.0/24 is subnetted, 3 subnets
C       150.1.23.0 is directly connected, FastEthernet0/0
C       150.1.3.0 is directly connected, Loopback1
B       150.1.2.0 [20/0] via 10.1.13.1, 00:15:22
```

Task 5

After implementing the previous task, the administrators realized that the traffic between networks 150.1.2.0 /24 and 150.1.3.0 /24 is taking a sub-optimal path; it is not using the direct path between routers R2 and R3.
Implement a BGP solution to fix this problem; you should NOT use the distance or any global config mode command to accomplish this task.

On R2

R2(config)#router bgp 200
R2(config-router)#**network 150.1.3.0 mask 255.255.255.0** backdoor

On R3

R3(config)#router bgp 300
R3(config-router)#**network 150.1.2.0 mask 255.255.255.0** backdoor

To verify the configuration:

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [20/0] via 10.1.12.1, 00:27:57
C    2.0.0.0/8 is directly connected, Loopback0
B    3.0.0.0/8 [20/0] via 10.1.12.1, 00:27:57
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.13.0 [120/1] via 10.1.12.1, 00:00:14, Serial0/0.21
C       10.1.12.0 is directly connected, Serial0/0.21
     150.1.0.0/24 is subnetted, 3 subnets
C       150.1.23.0 is directly connected, FastEthernet0/0
D       150.1.3.0 [90/156160] via 150.1.23.3, 00:01:19, FastEthernet0/0
C       150.1.2.0 is directly connected, Loopback1
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [20/0] via 10.1.13.1, 00:28:07
B    2.0.0.0/8 [20/0] via 10.1.13.1, 00:28:38
C    3.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
R       10.1.12.0 [120/1] via 10.1.13.1, 00:00:14, Serial0/0.31
     150.1.0.0/24 is subnetted, 3 subnets
```

```
C       150.1.23.0 is directly connected, FastEthernet0/0
C       150.1.3.0 is directly connected, Loopback1
D       150.1.2.0 [90/156160] via 150.1.23.2, 00:01:11, FastEthernet0/0
```

Note R2 and R3 were receiving routing information for networks 150.1.2.0 /24 and 150.1.3.0 /24 from two different sources, BGP and Eigrp.

R2 and R3 were using the routing information from BGP because it had a lower administrative distance (20 versus 90).
The "Network backdoor" command is a BGP solution to this problem; the BGP "backdoor" option assigns an administrative distance of 200 to networks 150.1.2.0 /24 and 150.1.3.0 /24, therefore, making the Eigrp more believable.

Enter the following commands to actually see the changed administrative distance:

On R2 and R3

```
(config)#NO router eigrp 100
```

On R2

R2#Sh ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [20/0] via 10.1.12.1, 00:36:39
C    2.0.0.0/8 is directly connected, Loopback0
B    3.0.0.0/8 [20/0] via 10.1.12.1, 00:36:39
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.13.0 [120/1] via 10.1.12.1, 00:00:11, Serial0/0.21
C       10.1.12.0 is directly connected, Serial0/0.21
     150.1.0.0/24 is subnetted, 3 subnets
C       150.1.23.0 is directly connected, FastEthernet0/0
B       150.1.3.0 [200/0] via 10.1.12.1, 00:00:13
C       150.1.2.0 is directly connected, Loopback1
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    1.0.0.0/8 [20/0] via 10.1.13.1, 00:39:19
B    2.0.0.0/8 [20/0] via 10.1.13.1, 00:39:50
C    3.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
R       10.1.12.0 [120/1] via 10.1.13.1, 00:00:19, Serial0/0.31
     150.1.0.0/24 is subnetted, 3 subnets
C       150.1.23.0 is directly connected, FastEthernet0/0
C       150.1.3.0 is directly connected, Loopback1
B       150.1.2.0 [200/0] via 10.1.13.1, 00:02:57
```

Task 6

Remove the IP address from the F0/0 interface of R2 and R3 and ensure that the F0/0 interface of both routers is in administratively down state. You should also remove the Loopback1 interface from these two routers.

On R2 and R3

```
(config)#Interface F0/0
(config)#NO ip address
(config)#Shutdown

(config)#NO int lo1
```

Task 7

Configure R1 as follows:

If network 2.0.0.0 is up and it's advertised to R1, R1 should take the following actions:

- R1 should NOT advertise it's network 1.0.0.0 /8 to R3.
- R1 should ONLY advertise network 2.0.0.0 /8 to R3

However, if network 2.0.0.0 /8 is down, then R1 should take the following actions:

- R1 should advertise network 1.0.0.0 /8 to R3.
- R1 should remove network 2.0.0.0 /8 from it's BGP table.

Before configuring this task you should verify the current BGP table of these routers:

R1#Show ip bgp

```
BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop       Metric  LocPrf    Weight    Path
*> 1.0.0.0          0.0.0.0             0              32768     i
*> 2.0.0.0          10.1.12.2           0                  0    200 i
*> 3.0.0.0          10.1.13.3           0                  0    300 i
```

R2#Show ip bgp

```
BGP table version is 7, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop       Metric  LocPrf    Weight    Path
*> 1.0.0.0          10.1.12.1           0                  0    100 i
*> 2.0.0.0          0.0.0.0             0              32768     i
*> 3.0.0.0          10.1.12.1                              0    100 300 i
```

R3#Show ip bgp

```
BGP table version is 7, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
```

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	10.1.13.1	0		0	100 i
*> 2.0.0.0	10.1.13.1			0	100 200 i
*> 3.0.0.0	0.0.0.0	0		32768	i

To conditionally advertise selected routes we can use the following commands:

- **Advertise-map**
- **Non-exist-map**
- **Exist-map**

This situation calls for the use of the "advertise-map" and non-exist-map" as follows:

On R1

R1(config)#access-list **1** permit 1.0.0.0 0.255.255.255

R1(config)#access-list **2** permit 2.0.0.0 0.255.255.255

R1(config)#route-map **ADV** permit 10
R1(config-route-map)#match ip addr **1**
R1(config-route-map)#exit

R1(config)#route-map **NotThere** permit 10
R1(config-route-map)#match ip addr **2**
R1(config-route-map)#exit

To prevent confusion you should select meaningful names for the route-maps. Note the access-list numbers and the names of the route-map.

R1(config)#router bgp 100
R1(config-router)#**neighbor 10.1.13.3 advertise-map ADV non-exist-map NotThere**

The neighbor command has the following route-maps:

- **The advertise-map – Specifies the name of the route-map that will be advertised if the condition of the non-exist-map is met.**

- **Non-exist-map – specifies the name of the route-map that will be compared to the advertise-map. If the condition is met and no match occurs, the route will be advertised. If a match occurs, then the condition is NOT met, and the route**

is withdrawn.

Note if network 2.0.0.0 is up, then network 1.0.0.0 should NOT be advertised to R3, since all the networks are up and advertised, R1 should withdraw it's network (1.0.0.0 /8):

On R1

R1#Show ip bgp

```
BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i
```

Note R1 does NOT advertise it's network (1.0.0.0 /8) to R3:

R1#Show ip bgp neighbors 10.1.13.3 advertised-routes

```
BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Note the output of the following command reveals that the bgp table of R3 does not have network 1.0.0.0 /8:

On R3

R3#Show ip bgp

```
BGP table version is 34, local router ID is 3.3.3.3

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the condition:

On R2

```
R2(config)#int lo0
R2(config-if)#Shut
```

The output of the following "Show" command reveals that network 2.0.0.0 is DOWN; therefore, R1 should advertise its network (1.0.0.0 /8) to R3. It may take few seconds for this policy to get implemented:

On R1

```
R1#Show ip bgp neighbors 10.1.13.3 advertised-routes

BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```

Note network 1.0.0.0 is advertised to R3.

Task 8

Remove the configuration commands entered in Task 7 before you proceed to the next task. Ensure that the routers have the advertised networks in their BGP table.

On R1

```
R1(config)#NO access-list 1
R1(config)#NO access-list 2

R1(config)#NO route-map ADV
```

R1(config)#NO route-map NotThere

R1(config)#router bgp 100
R1(config-router)#NO neighbor 10.1.13.3 advertise-map ADV non-exist-map NotThere

R1#Clear ip bgp *

On R2

R2(config)#int lo0
R2(config-if)#No shut

On R1

R1#Show ip bgp

```
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i
```

On R2

R2#Show ip bgp

```
BGP table version is 18, local router ID is 150.1.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.12.1                              0 100 300 i
```

On R3

R3#Show ip bgp

```
BGP table version is 19, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric  LocPrf  Weight  Path
*> 1.0.0.0          10.1.13.1          0                0 100 i
*> 2.0.0.0          10.1.13.1                           0 100 200 i
*> 3.0.0.0          0.0.0.0            0            32768  i
```

Task 9

R1 should be configured according to the following policy:

- If both networks (1.0.0.0 /8 and 2.0.0.0 /8) are up, then both networks should be advertised to R3.
- If network 1.0.0.0 /8 is down, R1 should NOT advertise network 2.0.0.0 /8 to R3.
- If network 2.0.0.0 /8 is down, then R1 should only advertise network 1.0.0.0 /8 to R3.

On R1

```
R1(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R1(config)#access-list 2 permit 2.0.0.0 0.255.255.255

R1(config)#route-map ADV permit 10
R1(config-route-map)#match ip addr 2
R1(config-route-map)#exit

R1(config)#route-map EXIST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#exit

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.13.3 advertise-map ADV exist-map EXIST
```

To test the first condition:

On R1

R1#Show ip bgp

```
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.13.3                0             0 300 i
```

On R2

R2#Show ip bgp

```
BGP table version is 18, local router ID is 150.1.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.12.1                              0 100 300 i
```

On R3

R3#Show ip bgp

```
BGP table version is 19, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.13.1                0             0 100 i
*> 2.0.0.0          10.1.13.1                              0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the second condition:

On R1

R1(config)#Int lo0
R1(config-if)#Shut

To test and verify the configuration:

On R1

R1#Show ip bgp neighbors 10.1.13.3 advertised-routes

Total number of prefixes 0

On R3

R3#Sh ip bgp

```
BGP table version is 12, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Note if network 1.0.0.0 is down none of the networks are advertised to R3.

To bring up the Loopback 0 interface of R1:

On R1

R1(config)#Int lo0
R1(config-if)#No Shut

To verify the configuration:

On R3

R3#Sh ip bgp

```
BGP table version is 14, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric  LocPrf  Weight  Path
*> 1.0.0.0          10.1.13.1            0               0 100 i
*> 2.0.0.0          10.1.13.1                            0 100 200 i
*> 3.0.0.0           0.0.0.0             0           32768  i
```

To test the third condition:

On R2

```
R2(config)#int lo0
R2(config-if)#Shut
```

On R1

R1#Show ip bgp neighbors 10.1.13.3 advertised-routes

```
BGP table version is 6, local router ID is 10.1.13.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf  Weight  Path
*> 1.0.0.0          0.0.0.0              0          32768   i

Total number of prefixes 1
```

On R3

R3#Show ip bgp

```
BGP table version is 17, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf  Weight  Path
*> 1.0.0.0          10.1.13.1            0               0 100 i
*> 3.0.0.0           0.0.0.0             0           32768  i
```

Task 10

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 4
Route Dampening

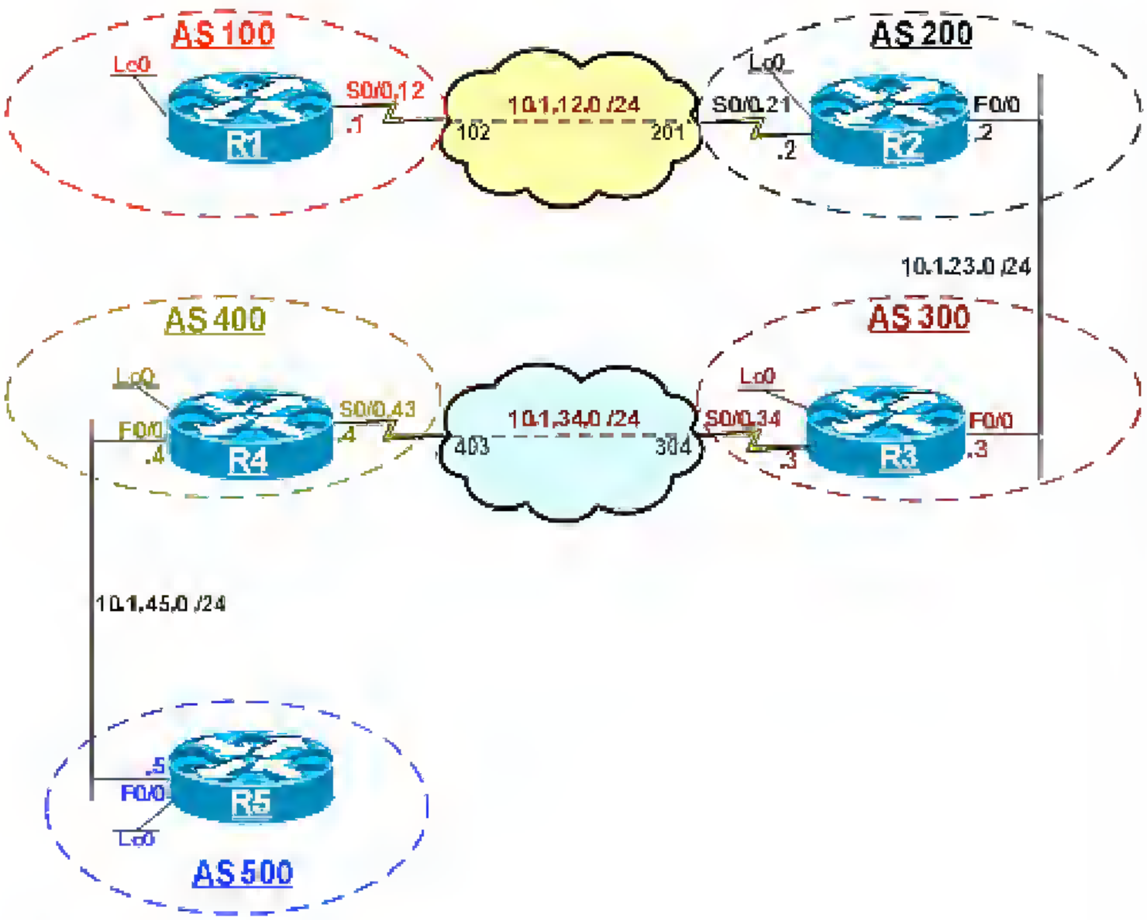

Lab Setup:

- Configure the frame-relay connections between the routers in a point-to-point manner.
- Configure R2 and R3's F0/0 interface in VLAN 23.
- Configure R4 and R5's F0/0 interface in VLAN 45.
- The IP address assignment of the routers should be based on the following IP addressing chart:

Ip addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	100
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 S0/0.34 F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24	300
R4	Lo0 Lo1 Lo2 S0/0.43 F0/0	4.4.4.4 /8 40.4.4.4 /24 44.4.4.4 /24 10.1.34.4 /24 10.1.45.4 /24	400
R5	Lo0 F0/0	5.5.5.5 /8 10.1.45.5 /24	500

Task 1

Configure an EBGP peer session between R1 and R2 and only advertise their Loopback0 interface in BGP. Ensure that these routers have NLRI to each others advertised prefix. R1 should be in AS 100 and R2 should be in AS 200.

On R1

```
R1(config)#router bgp 100
```

```
R1(config-router)#no au
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#network 1.0.0.0
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#network 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#no au
```

To verify the configuration:

```
R1#Show ip bgp

BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 2

Configure route dampening on R1 using the default parameters.

On R1

```
R1(config-router)#router bgp 100
R1(config-router)#bgp dampening
```

The parameters of BGP dampening are as follows:

- **Half-time** – Once a route has been assigned a penalty, the accumulated penalty is decreased every 5 seconds such that when the half period expires, the accumulated penalty is reduced by half. The default value of half-time is 15 minutes and the range is 1 to 45 minutes.

- **Reuse** – If the penalties for a flapping route is decreased enough to fall below

this value, the route is reusable. The default is 750 and the range is 1 to 20000.

- **Suppress** – Once the accumulated penalties reach this value, the route is suppressed. The default value is 2000 and the range is 1 – 20000.

- **Max-Suppress-Time** – The maximum time in minutes that a route can be suppressed. The default value is 4 times the half-time value (60 minutes) and the range is 1 to 255.

Therefore this configuration performs the following:

Half-time = 15 minutes, reuse = 750, Suppress = 2000 and Max-Suppress-Time = 60.

To see the parameters for dampening:

On R1

R1#Sh ip bgp dampening parameters

```
 dampening 15 750 2000 60 (DEFAULT)
  Half-life time          : 15 mins       Decay Time          : 2320 secs
  Max suppress penalty: 12000         Max suppress time: 60 mins
  Suppress penalty        :  2000          Reuse penalty        : 750
```

If network 2.0.0.0 is shutdown and then brought backup few times, the flap-statistics can be viewed in the "Show ip bgp dampening flap-statistics" command.

Task 3

Configure an EBGP peer session between R2 and R3, and advertise their Loopback0 interface in BGP. Ensure that these routers have NLRI to each others Loopback interface. R3 should be configured in AS 300.

On R2

```
R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3

```
R3(config)#router bgp 300
```

```
R3(config-router)#no au
R3(config-router)#network 3.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
```

To verify the configuration:

On R3

R3#Show ip bgp

```
BGP table version is 4, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric  LocPrf   Weight   Path
*> 1.0.0.0          10.1.23.2                                0    200 100 i
*> 2.0.0.0          10.1.23.2                0               0    200 i
*> 3.0.0.0          0.0.0.0                  0           32768    i
```

Task 4

Configure route dampening on R3 such that the half-life parameter is set to 30 minutes. This router should use the default parameters for suppress-limit, reuse, and maximum suppress time.

On R3

R3(config-router)#**bgp dampening 30 750 2000 60**

To verify the configuration:

On R3

R3#Sh ip bgp dampening parameters

```
dampening 30 750 2000 60
 Half-life time         : 30 mins      Decay Time          : 1045 secs
 Max suppress penalty:  3000           Max suppress time: 60 mins
 Suppress penalty       :  2000        Reuse penalty        : 750
```

Note you may get a "% dampening reconfiguration in progress for IPv4 Unicast" message, if you do, you should wait few seconds and try again.

Task 5

Configure an EBGP session between R3 and R4. Advertise Loopback0, Loopback1 and Loopback2 interface of R4 in BGP. Router R4 should be configured in AS 400.

On R4

```
R4(config)#router bgp 400
R4(config-router)#netw 4.0.0.0
R4(config-router)#netw 40.4.4.0 mask 255.255.255.0
R4(config-router)#netw 44.4.4.0 mask 255.255.255.0
R4(config-router)#neighbor 10.1.34.3 remote-as 300
R4(config-router)#no au
```

On R3

```
R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.34.4 remote-as 400
```

To verify the configuration:

On R4

R4#Show ip bgp

```
BGP table version is 7, local router ID is 44.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 3.0.0.0          10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 40.4.4.0/24      0.0.0.0                  0         32768 i
*> 44.4.4.0/24      0.0.0.0                  0         32768 i
```

Task 6

Configure an EBGP peer session between R4 and R5. R5 should advertise its Loopback 0 interface in BGP. Ensure that these routers have NLRI to each others Loopback interface. R5 should be configured in AS 500.

On R5

```
R5(config)#router bgp 500
R5(config-router)#no au
R5(config-router)#neighbor 10.1.45.4 remote-as 400
R5(config-router)#netw 5.0.0.0
```

On R4

```
R4(config)#router bgp 400
R4(config-router)#neighbor 10.1.45.5 remote-as 500
```

To verify the configuration:

R5#Show ip bgp

```
BGP table version is 8, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 400 300 200 100 i
*> 2.0.0.0          10.1.45.4                              0 400 300 200 i
*> 3.0.0.0          10.1.45.4                              0 400 300 i
*> 4.0.0.0          10.1.45.4                0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 40.4.4.0/24      10.1.45.4                0             0 400 i
*> 44.4.4.0/24      10.1.45.4                0             0 400 i
```

Task 7

Configure route dampening on R4 as follows:

- Network 40.4.4.0 /24 should have the following dampening parameters applied: **Max-Suppress-Time** of 90, **Reuse** 800, **Suppress** 2400 and a **Half-Time** of 20

➢ Network 44.4.4.0 /24 should have the following dampening parameters applied: **Max-Suppress-Time** of 60, **Reuse** 700, **Suppress** 2000 and a **Half-Time** of 15

On R4

R4(config)#access-list 40 permit 40.4.4.0 0.0.0.255

R4(config)#access-list 44 permit 44.4.4.0 0.0.0.255

R4(config)#route-map TST permit 10
R4(config-route-map)#match ip addr 40
R4(config-route-map)#set dampening 20 800 2400 90

R4(config)#route-map TST permit 20
R4(config-route-map)#match ip addr 44
R4(config-route-map)#set dampening 15 700 2000 60

R4(config)#route-map TST permit 30

R4(config)#router bgp 400
R4(config-router)#bgp dampening route-map TST

Note the route-map gives us flexibility. In this case we have applied two different route dampening parameters to different routes. The two networks are identified with access-lists. The route-map references the access-lists and sets the dampening parameters based on the networks.

To verify the configuration:

On R4

R4#Sh ip bgp damp parameters

```
dampening 20 800 2400 90 (route-map TST 10)
 Half-life time         : 20 mins     Decay Time        : 3490 secs
 Max suppress penalty: 18075          Max suppress time: 90 mins
 Suppress penalty       :  2400        Reuse penalty       : 800

dampening 15 700 2000 60 (route-map TST 20)

 Half-life time         : 15 mins     Decay Time        : 2235 secs
 Max suppress penalty: 11200          Max suppress time: 60 mins
 Suppress penalty       :  2000       Reuse penalty       : 700
```

Task 8

Configure route dampening on R2 using the following policy:

- All the existing and future prefixes from AS 300 should have the following parameters applied:

  Max-Suppress-Time of 80
  Reuse 750
  Suppress 2200
  Half-Time 30

On R2

```
R2(config)#ip as-path access-list 1 permit ^300$

R2(config)#route-map TST permit 10
R2(config-route-map)#match as-path 1
R2(config-route-map)#set dampening 30 750 2200 80

R2(config)#route-map TST permit 20

R2(config)#router bgp 200
R2(config-router)#bgp dampening route-map TST
```

The combination of "route-map" and the "as-path access-list" command can apply bgp dampening to an AS based on the AS number.

To verify the configuration:

On R2

R2#Show ip bgp damp parameters

```
 dampening 30 750 2200 80 (route-map TST 10)
  Half-life time       : 30 mins       Decay Time       : 1995 secs
  Max suppress penalty:  4755          Max suppress time: 80 mins

  Suppress penalty     :  2200         Reuse penalty     : 750
```

Task 9

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 5
Route Aggregation

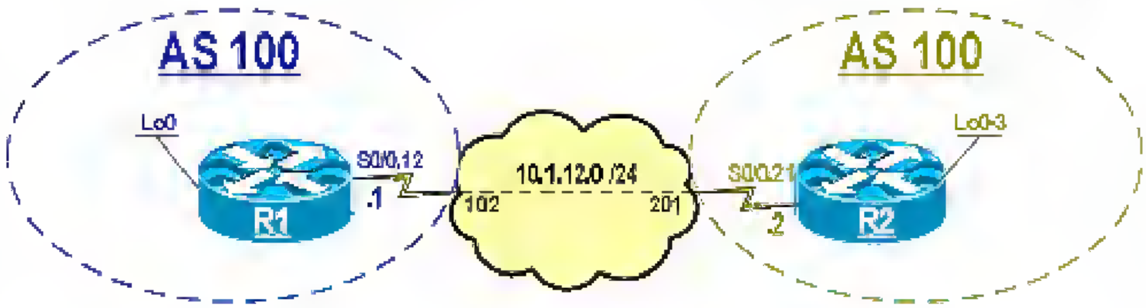

Lab Setup:

- Configure frame-relay connection between the routers in a point-to-point manner.
- Use the following IP addressing chart for IP assignment.

Ip addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Frame-relay connection to R2	1.1.1.1 /8 10.1.12.1 /24	100
R2	Lo0 Lo1 Lo2 Lo3 Frame-relay connection to R1	2.2.0.2 / 24 2.2.1.2 /24 2.2.2.2 /24 2.2.3.2 /24 10.1.12.2 /24	200

Task 1

Configure an EBGP session between the routers and only advertise their Loopback interface/s in BGP. R1 should be in AS 100 and R2 should be configured in AS 200.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#network 2.2.0.0 mask 255.255.255.0
R2(config-router)#network 2.2.1.0 mask 255.255.255.0
R2(config-router)#network 2.2.2.0 mask 255.255.255.0
R2(config-router)#network 2.2.3.0 mask 255.255.255.0
```

To verify the configuration

R2#Show ip bgp

```
BGP table version is 6, local router ID is 2.2.3.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.2.0.0/24       0.0.0.0                  0         32768 i
*> 2.2.1.0/24       0.0.0.0                  0         32768 i
*> 2.2.2.0/24       0.0.0.0                  0         32768 i
*> 2.2.3.0/24       0.0.0.0                  0         32768 i
```

Task 2

Configure R2 such that it summarizes it's Loopback interfaces and advertises a single summary to R1. R2 should NOT assign an atomic-aggregate to the summary route when it advertises it to any of its neighbors.

Note in BGP, an aggregate is only created if at least one of the specific routes of the aggregate exists in the BGP table. It is recommended to configure most if not all the specific routes with a Network statement, because if only a single Network is configured to satisfy the requirements, and that particular Network goes down, then the aggregate

will be removed.
There are many ways to advertise an aggregate, one way to advertise an aggregate is by creating a static route that matches the aggregate route and then advertising the aggregate in BGP as follows:

```
(config)#Ip route 2.2.0.0 255.255.252.0 null 0

(config)#Router bgp 200
(config-router)#Network 2.2.0.0 mask 255.255.252.0
```

But since that is not an option here, we had to advertise every specific prefix under the aggregate.

On R2

```
R2(config)#Router bgp 200
R2(config-router)#aggregate-address 2.2.0.0 255.255.252.0
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.2.0.0/24       10.1.12.2                0             0 200 i
*> 2.2.0.0/22       10.1.12.2                0             0 200 i
*> 2.2.1.0/24       10.1.12.2                0             0 200 i
*> 2.2.2.0/24       10.1.12.2                0             0 200 i
*> 2.2.3.0/24       10.1.12.2                0             0 200 i
```

On R2

R2#Show ip bgp

```
BGP table version is 7, local router ID is 2.2.3.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1            0             0    100 i
*> 2.2.0.0/24       0.0.0.0              0               32768 i
*> 2.2.0.0/22       0.0.0.0                              32768 i
*> 2.2.1.0/24       0.0.0.0              0               32768 i
*> 2.2.2.0/24       0.0.0.0              0               32768 i
*> 2.2.3.0/24       0.0.0.0              0               32768 i
```

By default in BGP the aggregate and all the specific routes are advertised. A "summary-only" argument used with the aggregate-address will suppress the specific routes so that only the aggregate route is advertised.

Note none of the prefixes are suppresses.

On R2

```
R2(config)#router bgp 200
R2(config-router)#aggregate-address 2.2.0.0 255.255.252.0 summary-only
```

Note the "summary-only" keyword will suppress all the prefixes on R2 such that R2's neighbors won't see the more specific routes.

On R2

R2#Show ip bgp

```
BGP table version is 11, local router ID is 2.2.3.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1            0             0 100 i
s> 2.2.0.0/24       0.0.0.0              0         32768 i
*> 2.2.0.0/22       0.0.0.0                        32768 i
s> 2.2.1.0/24       0.0.0.0              0         32768 i
s> 2.2.2.0/24       0.0.0.0              0         32768 i
s> 2.2.3.0/24       0.0.0.0              0         32768 i
```

Note the letter "s" to the left of the ">" sign. The "s" means that these prefixes are suppressed.

On R1

R1#Show ip bgp

```
BGP table version is 11, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.2.0.0/22       10.1.12.2                0             0 200 i
```

Note the only route that is advertised to R1 is the summary route.

On R1

R1#Sh ip bgp 2.2.0.0/22

```
BGP routing table entry for 2.2.0.0/22, version 17
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200, (aggregated by 200 2.2.3.2)
    10.1.12.2 from 10.1.12.2 (2.2.3.2)
      Origin IGP, metric 0, localpref 100, valid, external, atomic-aggregate, best
```

On R2

R2#Show ip bgp 2.2.0.0/22

```
BGP routing table entry for 2.2.0.0/22, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     1
  Local, (aggregated by 200 2.2.3.2)
    0.0.0.0 from 0.0.0.0 (2.2.3.2)
      Origin IGP, localpref 100, weight 32768, valid, aggregated, local, atomic-aggregate, best
```

Note the output of the "Show ip bgp 2.2.0.0/22" command above, displays two different attributes, the "aggregator" and the "atomic-aggregate" attribute.

The "aggregator" attribute identifies the AS number that the aggregation was performed and it also identifies the router-id of the router that performed the aggregation.

By default when aggregation is configured in BGP, the "atomic-aggregate" attribute is attached to the aggregate address; this alarms the administrator that certain information could be hidden.

We know that the specific routes under that aggregate are always suppressed/hidden when the summary or aggregation is performed under any routing protocol, but in BGP another hidden or suppressed item is the actual AS number/s that the specific routes were originated.

An atomic-aggregate – This is an attribute that is assigned to the aggregate route automatically if the "as-set" argument is not used in the "aggregate-address" command. When an aggregation is performed, certain information is lost. In BGP that information is not only the more specific routes under that aggregate, but it can also be the AS numbers that the prefixes traversed through to get to the router that is performing the aggregation. If it's not corrected a routing loop can occur. In order to prevent the routing loops from occurring, the "AS-SET" argument should be used when performing aggregation.

The "AS-SET" argument used in the aggregate-address command reveals the AS number/s that some, if not all the specific routes were originated from, once that information is revealed, the "atomic-aggregate" attribute is automatically removed.

On R2

R2(config)#router bgp 200
R2(config-router)#**aggregate-address 2.2.0.0 255.255.252.0 summary-only** as-set

To verify the configuration:

On R1

R1#Sh ip bgp 2.2.0.0/22

```
BGP routing table entry for 2.2.0.0/22, version 22
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200, (aggregated by 200 2.2.3.2)
    10.1.12.2 from 10.1.12.2 (2.2.3.2)
      Origin IGP, metric 0, localpref 100, valid, external, best
```

On R2

R2#Sh ip bgp 2.2.0.0/22

```
BGP routing table entry for 2.2.0.0/22, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)

  Advertised to update-groups:
     1
  Local, (aggregated by 200 2.2.3.2)
    0.0.0.0 from 0.0.0.0 (2.2.3.2)
      Origin IGP, localpref 100, weight 32768, valid, aggregated, local, best
```

Note the atomic-aggregate is no longer attached to the aggregate-address.

Task 3

Reconfigure the routers using the follows diagram/IP addressing information and ONLY advertise their Loopback interfaces in BGP. You can use the initial config file for advertising and setting up the diagram.

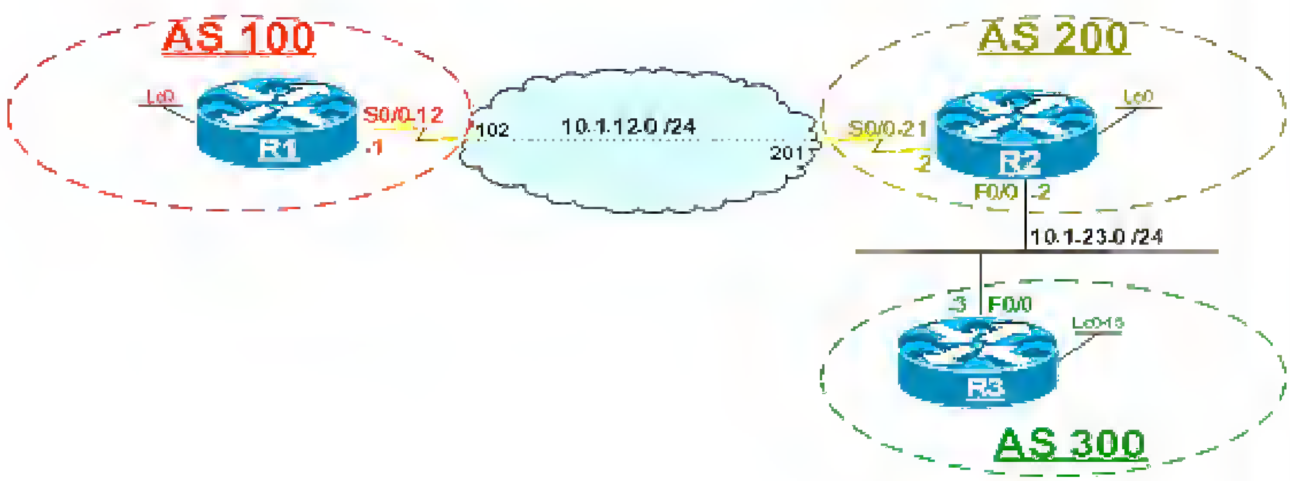

Lab Setup:

- Configure a frame-relay point-to-point connection between routers R1 and R2.
- Configure the F0/0 interface of R2 and R3 in VLAN 23.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	3.1.11.1 /24 10.1.12.1 /24	100
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**200**
R3	Lo0 Lo1 Lo2 Lo3 Lo4 Lo5 Lo6 Lo7 Lo8 Lo9 Lo10 Lo12 Lo13 Lo14 Lo15 F0/0	3.1.0.3 /24 3.1.1.3 /24 3.1.2.3 /24 3.1.3.3 /24 3.1.4.3 /24 3.1.5.3 /24 3.1.6.3 /24 3.1.7.3 /24 3.1.8.3 /24 3.1.9.3 /24 3.1.10.3 /24 3.1.12.3 /24 3.1.13.3 /24 3.1.14.3 /24 3.1.15.3 /24 10.1.23.3 /24	300

Task 4

Configure router R1 in AS 100 to establish an EBGP session with R2 in AS 200, and router R2 in AS 200 should establish an EBGP peer session with R3 in AS 300. These routers should advertise their loopback interface/s in their AS.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no au
```

R1(config-router)#netw 3.1.11.0 mask 255.255.255.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200

On R2

R2(config)#router bgp 200
R2(config-router)#no au
R2(config-router)#neighbor 10.1.23.3 remote-as 300
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#network 2.0.0.0

On R3

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#netw 3.1.0.0 mask 255.255.255.0
R3(config-router)#netw 3.1.1.0 mask 255.255.255.0
R3(config-router)#netw 3.1.2.0 mask 255.255.255.0
R3(config-router)#netw 3.1.3.0 mask 255.255.255.0
R3(config-router)#netw 3.1.4.0 mask 255.255.255.0
R3(config-router)#netw 3.1.5.0 mask 255.255.255.0
R3(config-router)#netw 3.1.6.0 mask 255.255.255.0
R3(config-router)#netw 3.1.7.0 mask 255.255.255.0
R3(config-router)#netw 3.1.8.0 mask 255.255.255.0
R3(config-router)#netw 3.1.9.0 mask 255.255.255.0
R3(config-router)#netw 3.1.10.0 mask 255.255.255.0
R3(config-router)#netw 3.1.12.0 mask 255.255.255.0
R3(config-router)#netw 3.1.13.0 mask 255.255.255.0
R3(config-router)#netw 3.1.14.0 mask 255.255.255.0
R3(config-router)#netw 3.1.15.0 mask 255.255.255.0

The configuration for advertising the 12 prefixes can be downloaded from the CD provided with this work book

To verify the configuration:

On R2

R2#Sh ip bgp

BGP table version is 15, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

```
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0              0         32768 i
*> 3.1.0.0/24       10.1.23.3            0             0 300 i
*> 3.1.1.0/24       10.1.23.3            0             0 300 i
*> 3.1.2.0/24       10.1.23.3            0             0 300 i
*> 3.1.3.0/24       10.1.23.3            0             0 300 i
*> 3.1.4.0/24       10.1.23.3            0             0 300 i
*> 3.1.5.0/24       10.1.23.3            0             0 300 i
*> 3.1.6.0/24       10.1.23.3            0             0 300 i
*> 3.1.7.0/24       10.1.23.3            0             0 300 i
*> 3.1.8.0/24       10.1.23.3            0             0 300 i
*> 3.1.9.0/24       10.1.23.3            0             0 300 i
*> 3.1.10.0/24      10.1.23.3            0             0 300 i
*> 3.1.11.0/24      10.1.12.1            0             0 100 i
*> 3.1.12.0/24      10.1.23.3            0             0 300 i
*> 3.1.13.0/24      10.1.23.3            0             0 300 i
*> 3.1.14.0/24      10.1.23.3            0             0 300 i
*> 3.1.15.0/24      10.1.23.3            0             0 300 i
```

Task 5

R2 should aggregate all the networks in 3.1.0.0 address space and advertise a single aggregate route that only aggregates the specific routes for subnets under the 3.0.0.0 network in it's BGP table, ensure that the atomic-aggregate attribute is not attached to the aggregate route. This aggregation should be configured such that R1 in AS 100 is the only AS that receives the aggregate route, R3 in AS 300 should NOT receive the aggregate route. R1 should use R2 as the next hop to reach any of the specific routes within the aggregate, R1 should NOT use R2 if it's network 3.1.11.0 /24 network is down. R3 does NOT need NLRI to network 3.1.11.0 /24 advertised by R1.

On R2

R2 should aggregate all the networks in 3.1.0.0 address space and advertise a single aggregate route that only aggregates the specific routes for subnets under the 3.0.0.0 network in it's BGP table, ensure that the atomic-aggregate attribute is not attached to the aggregate route.

R2(config)#Router bgp 200
R2(config-router)#**aggregate-address 3.1.0.0 255.255.240.0** summary-only as-set

This command aggregates networks 3.1.0.0 /24 – 3.1.15.0 /24 and only advertises the summary route and not all the specific routes, the "summary-only" argument accomplish that. This aggregate route will not have the "atomic-aggregate" attribute attached because the "as-set" argument is used.

To verify the configuration:

On R2

R2#Sh ip bgp

```
BGP table version is 63, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
s> 3.1.0.0/24       10.1.23.3                0             0 300 i
*> 3.1.0.0/20       0.0.0.0                       100  32768 {300,100} i
s> 3.1.1.0/24       10.1.23.3                0             0 300 i
s> 3.1.2.0/24       10.1.23.3                0             0 300 i
s> 3.1.3.0/24       10.1.23.3                0             0 300 i
s> 3.1.4.0/24       10.1.23.3                0             0 300 i
s> 3.1.5.0/24       10.1.23.3                0             0 300 i
s> 3.1.6.0/24       10.1.23.3                0             0 300 i
s> 3.1.7.0/24       10.1.23.3                0             0 300 i
s> 3.1.8.0/24       10.1.23.3                0             0 300 i
s> 3.1.9.0/24       10.1.23.3                0             0 300 i
s> 3.1.10.0/24      10.1.23.3                0             0 300 i
s> 3.1.11.0/24      10.1.12.1                0             0 100 i
s> 3.1.12.0/24      10.1.23.3                0             0 300 i
s> 3.1.13.0/24      10.1.23.3                0             0 300 i
s> 3.1.14.0/24      10.1.23.3                0             0 300 i
s> 3.1.15.0/24      10.1.23.3                0             0 300 i
```

Note the specific routes are suppressed (The letter "S" to the left of the ">" sign). The curly brackets in the path column identify the AS numbers that the specific routes originated from. Because the AS numbers are now included in the path column, neither R1 in AS 100 nor R3 in AS 300 will have the aggregate route in their BGP table.

On R1

R1#Sh ip bgp

```
BGP table version is 130, local router ID is 3.1.12.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 49, local router ID is 3.3.11.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.1.0.0/24       0.0.0.0                  0         32768 i
*> 3.1.1.0/24       0.0.0.0                  0         32768 i
*> 3.1.2.0/24       0.0.0.0                  0         32768 i
*> 3.1.3.0/24       0.0.0.0                  0         32768 i
*> 3.1.4.0/24       0.0.0.0                  0         32768 i
*> 3.1.5.0/24       0.0.0.0                  0         32768 i
*> 3.1.6.0/24       0.0.0.0                  0         32768 i
*> 3.1.7.0/24       0.0.0.0                  0         32768 i
*> 3.1.8.0/24       0.0.0.0                  0         32768 i
*> 3.1.9.0/24       0.0.0.0                  0         32768 i
*> 3.1.10.0/24      0.0.0.0                  0         32768 i
*> 3.1.12.0/24      0.0.0.0                  0         32768 i
*> 3.1.13.0/24      0.0.0.0                  0         32768 i
*> 3.1.14.0/24      0.0.0.0                  0         32768 i
*> 3.1.15.0/24      0.0.0.0                  0         32768 i
```

Note R1 and R3 do NOT have the aggregate route in their BGP table.

This aggregation should be configured such that R1 in AS 100 is the only AS that receives the aggregate route, R3 or future peer neighbors should NOT receive the aggregate route. R1 should use R2 as the next hop to reach any of the specific routes within the aggregate

R2(config)#**ip as-path access-list 1 permit ^300$**

The above command identifies AS number 300.

R2(config)#route-map TST permit 10
R2(config-route-map)#match as-path 1

The "route-map" command references the "as-path access-list 1".

R2(config-router)#aggregate-address 3.1.0.0 255.255.240.0 as-set summary-only advertise-map TST

The "advertise-map" command assigns the route-map "TST" to the "aggregate address" command.

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 12, local router ID is 3.1.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop      Metric LocPrf   Weight Path
*> 2.0.0.0          10.1.12.2         0                0 200 i
*> 3.1.0.0/20       10.1.12.2         0           0 200 300 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 49, local router ID is 3.3.11.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop      Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2         0              0 200 i
*> 3.1.0.0/24       0.0.0.0           0          32768 i
*> 3.1.1.0/24       0.0.0.0           0          32768 i
*> 3.1.2.0/24       0.0.0.0           0          32768 i
*> 3.1.3.0/24       0.0.0.0           0          32768 i
*> 3.1.4.0/24       0.0.0.0           0          32768 i
```

```
*> 3.1.5.0/24       0.0.0.0                  0              32768 i
*> 3.1.6.0/24       0.0.0.0                  0              32768 i
*> 3.1.7.0/24       0.0.0.0                  0              32768 i
*> 3.1.8.0/24       0.0.0.0                  0              32768 i
*> 3.1.9.0/24       0.0.0.0                  0              32768 i
*> 3.1.10.0/24      0.0.0.0                  0              32768 i
*> 3.1.12.0/24      0.0.0.0                  0              32768 i
*> 3.1.13.0/24      0.0.0.0                  0              32768 i
*> 3.1.14.0/24      0.0.0.0                  0              32768 i
*> 3.1.15.0/24      0.0.0.0                  0              32768 i
```

Note R3 gets the aggregate route but it rejects it because it sees its own AS number in the as-path list. R1 receives and processes the aggregate route because it does not see its own AS number in the as-path list advertised by R2.

The following shows all the routes that are advertised by R2 to its neighbor 10.1.23.3:

On R2

R2#Show ip bgp neighbors 10.1.23.3 advertised-routes

```
BGP table version is 64, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.1.0.0/20       0.0.0.0                       100  32768 300 i

Total number of prefixes 2
```

The output of the following display shows all the routes received and accepted by R3:

On R3

R3#Show ip bgp neighbor 10.1.23.2 routes

```
BGP table version is 49, local router ID is 3.3.11.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
```

```
*> 2.0.0.0          10.1.23.2             0                   0 200 i
```

R1 should NOT use R2 if it's network 3.1.11.0 /24 network is down

On R1

R1(config)#**ip route 3.1.11.0 255.255.255.0 NULL0**

To verify the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    2.0.0.0/8 [20/0] via 10.1.12.2, 01:55:29
     3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
B       3.1.0.0/20 [20/0] via 10.1.12.2, 00:18:34
C       3.1.11.0/24 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
```

Note 3.1.11.0 /24 is directly connected, to test this condition we should shut down the interface and check the routing table again, as follows:

On R1

R1(config)#int lo0
R1(config-if)#Shut

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B – BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

B    2.0.0.0/8 [20/0] via 10.1.12.2, 01:58:14
     3.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
B       3.1.0.0/20 [20/0] via 10.1.12.2, 00:21:19
S       3.1.11.0/24 is directly connected, Null0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0.12
```

Note if the interface is down all the traffic destined for the network is forwarded to the NULL0 interface, and therefore, it won't be forwarded to R2.

Task 6

Configure R2 such that a cost of 50 is assigned to the aggregate route.

On R2

```
R2(config)#route-map COST permit 10
R2(config-route-map)#set metric 50
R2(config)#route-map COST permit 20

R2(config)#router bgp 200
R2(config-router)# aggregate-address 3.1.0.0 255.255.240.0 summary-only advertise-map TST as-set attribute-map COST
```

The attribute map identifies the name of the route-map used to set the attribute/s of the aggregate route.

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 13, local router ID is 3.1.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop      Metric LocPrf   Weight Path
*> 2.0.0.0          10.1.12.2          0               0 200 i
*> 3.1.11.0/24      0.0.0.0            0           32768 i
*> 3.1.0.0/20       10.1.12.2         50               0 200 300 i
```

Note R1 get the aggregate route with a cost of 50.

Task 7

The policy for R1 requires that it should receive the aggregate route plus one of the more specific subnets (3.1.3.0 /24). This policy should be configured and tested in three different ways using a suppress-map, and unsuppress-map.

To test the Suppress-map scenario #1:

On R2

R2(config)#access-list 1 deny 3.1.3.0 0.0.0.255
R2(config)#access-list 1 permit any

R2(config)#route-map **SUPP** permit 10
R2(config-route-map)#match ip addr 1

R2(config)#router bgp 200
R2(config-router)# **aggregate-address 3.1.0.0 255.255.240.0 summary-only advertise-map TST as-set attribute-map COST suppress-map SUPP**

To verify the configuration:

On R1

R1#Show ip bgp

BGP table version is 14, local router ID is 3.1.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

```
         r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                              0 200 300 i
```

To test the suppress-map scenario #2:

On R2

```
R2(config)#NO access-list 1
R2(config)#access-list 1 permit 3.1.3.0 0.0.0.255

R2(config)#Route-map SUPP deny 10
R2(config-route-map)#Match ip addr 1
R2(config)#Route-map SUPP permit 20

R2(config)#router bgp 200
R2(config-router)#NO  aggregate-address 3.1.0.0 255.255.240.0 summary-only
advertise-map TST as-set attribute-map COST suppress-map SUPP
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 14, local router ID is 3.1.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
         r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.1.11.0/24      0.0.0.0                  0         32768 i
*> 3.1.0.0/20       10.1.12.2               50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                              0 200 300 i
```

To test Unsuppress-map scenario #3:

R2(config)#NO route-map SUPP

R2(config)#route-map SUPP permit 10
R2(config-route-map)#match ip addr 1

R2(config)#Router bgp 200
R2(config-router)#**aggregate-address 3.1.0.0 255.255.240.0** summary-only **advertise-map** TST **as-set attribute-map** COST

This command takes off the suppress-map

R2(config-router)#**neighbor 10.1.12.1 unsuppress-map** **SUPP**

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 14, local router ID is 3.1.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf   Weight Path
*> 2.0.0.0          10.1.12.2           0               0 200 i
*> 3.1.11.0/24      0.0.0.0             0           32768 i
*> 3.1.0.0/20       10.1.12.2          50               0 200 300 i
*> 3.1.3.0/24       10.1.12.2                           0 200 300 i
```

Task 8

Remove the configuration command/s from the previous step.

On R2

R2(config)#NO **route-map SUPP**

R2(config)#NO **access-list**

R2(config)#Router bgp 200
R2(config-router)#NO **neighbor 10.1.12.1 unsuppress-map** SUPP

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 14, local router ID is 3.1.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf     Weight Path
*> 2.0.0.0          10.1.12.2           0                  0 200 i
*> 3.1.11.0/24      0.0.0.0             0              32768 i
*> 3.1.0.0/20       10.1.12.2          50                  0 200 300 i
```

Task 9

Configure R1 so it has the aggregate route plus the specific route that it wanted to have in its BGP table (3.1.3.0 /24). R1 should NOT advertise this subnet, configure a static route or use the redistribute command to accomplish this task. R2 should NOT be configured for this task.

To accomplish this task on R1, we can use the combination of Exist-map and Inject –map. The Exist-map matches on the aggregate address and the router that advertised the aggregate address (route-source command in the route-map called "EXIST"). The Inject-map injects the IP address/es identified by the route-map called INJECT, if the condition of the Exist-map is true. Therefore, if the router that advertised the aggregate address and the aggregate address exist, then, inject what ever that is specified in the prefix-list that is referenced by the route-map INJECT.

To verify the BGP table of R1 before the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 14, local router ID is 3.1.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop         Metric LocPrf      Weight Path
*> 2.0.0.0          10.1.12.2             0                  0 200 i
*> 3.1.11.0/24      0.0.0.0               0              32768 i
*> 3.1.0.0/20       10.1.12.2            50                  0 200 300 i
```

To configure:

On R1

```
R1(config)#ip prefix-list NET permit 3.1.3.0/24
R1(config)#ip prefix-list AGG permit 3.1.0.0/20
R1(config)#ip prefix-list R2 permit 10.1.12.2/32

R1(config)#route-map EXIST permit 10
R1(config-route-map)#match ip addr prefix-list AGG
R1(config-route-map)#match ip route-source prefix-list R2

R1(config)#route-map INJECT permit 10
R1(config-route-map)#set ip addr prefix-list NET

R1(config)#Router bgp 100
R1(config-router)#bgp inject-map INJECT exist-map EXIST
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 6, local router ID is 10.1.12.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop         Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2             0             0 200 i
*> 3.1.0.0/20       10.1.12.2            50             0 200 300 i
*> 3.1.3.0/24       10.1.12.2                           0 ?
*> 3.1.11.0/24      0.0.0.0               0         32768 i
```

Note the subnet is in the BGP table of R1.

Task 10

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 6
The Community Attribute

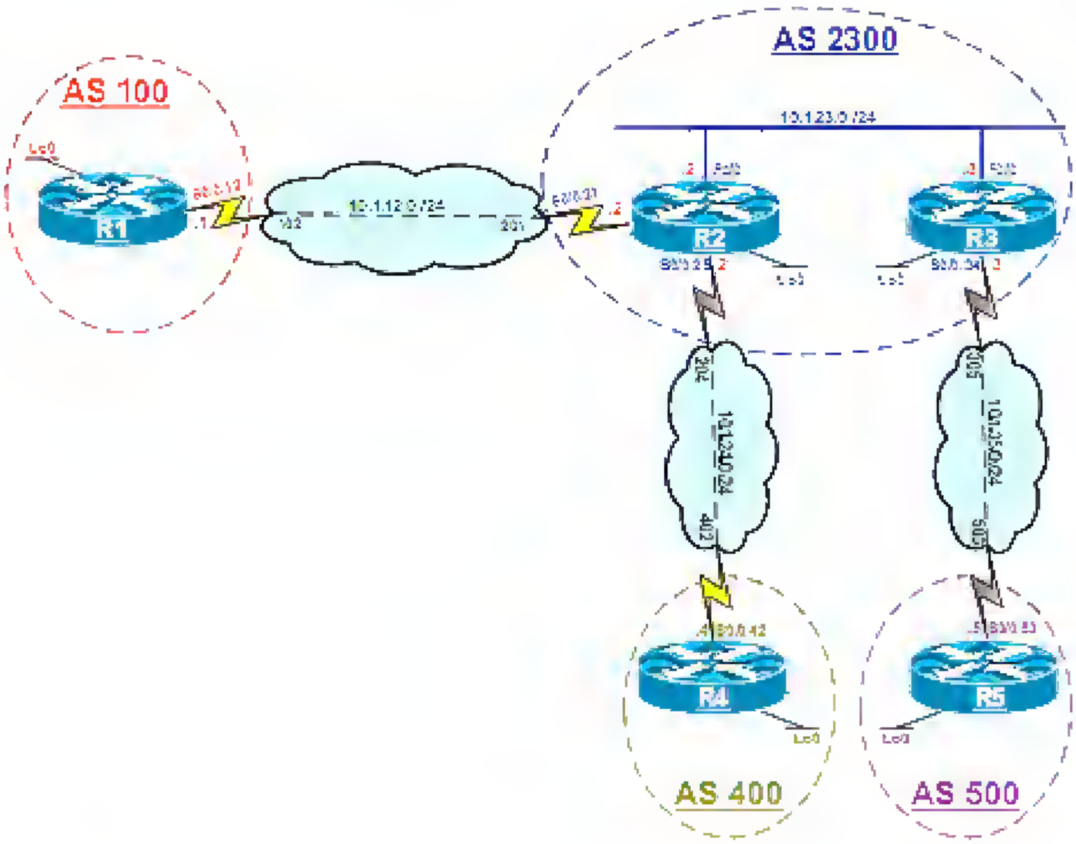

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP address chart for IP assignment.

AS & IP addressing chart:

Router	Interface	IP Address	AS Number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	100
R2	**Lo0** **S0/0.21** **S0/0.24** **F0/0 interface connection to R3**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.24.2 /24** **10.1.23.2 /24**	**2300**
R3	Lo0 F0/0 interface connection to R2 S0/0.35	3.3.3.3 /8 10.1.23.3 /24 10.1.35.3 /24	2300
R4	Lo0 **S0/0.42**	4.4.4.4 /8 **10.1.24.4 /24**	400
R5	Lo0 S0/0.53	5.5.5.5 /8 10.1.35.5 /24	500

Task 1

Configure EBGP peer session/s between the routers based on the above "AS & IP addressing chart". These routers should ONLY advertise their Loopback interface/s in BGP. These BGP routers should use their Loopback0's IP address as their Router id. Ensure that every router has NLRI to every link in this lab using RIPv2.

On All Routers

```
(config)#Router rip
(config-router)#No au
(config-router)#Ver 2
(config-router)#Network 10.0.0.0
```

On R1

```
R1(config)#router bgp 100
R1(config-router)#bgp router-id 1.1.1.1
R1(config-router)#netw 1.0.0.0
R1(config-router)#no au
R1(config-router)#neighbor 10.1.12.2 remote-as 2300
```

On R2

```
R2(config)#router bgp 2300
```

R2(config-router)#bgp router-id 2.2.2.2
R2(config-router)#netw 2.0.0.0
R2(config-router)#no au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 2300
R2(config-router)#neighbor 10.1.24.4 remote-as 400

On R3

R3(config)#router bgp 2300
R3(config-router)#bgp router-id 3.3.3.3
R3(config-router)#netw 3.0.0.0
R3(config-router)#no au
R3(config-router)#neighbor 10.1.23.2 remote-as 2300
R3(config-router)#neighbor 10.1.35.5 remote-as 500

On R4

R4(config)#router bgp 400
R4(config-router)#bgp router-id 4.4.4.4
R4(config-router)#netw 4.0.0.0
R4(config-router)#no au
R4(config-router)#neighbor 10.1.24.2 remote-as 2300

On R5

R5(config)#router bgp 500
R5(config-router)#bgp router-id 5.5.5.5
R5(config-router)#netw 5.0.0.0
R5(config-router)#no au
R5(config-router)#neighbor 10.1.35.3 remote-as 2300

To verify the configuration:

On R1

R1#Show ip bgp

BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
    Network          Next Hop        Metric LocPrf    Weight Path
*> 1.0.0.0          0.0.0.0              0                32768 i
*> 2.0.0.0          10.1.12.2            0                    0 2300 i
*> 3.0.0.0          10.1.12.2                                 0 2300 i
*> 4.0.0.0          10.1.12.2                                 0 2300 400 i
*> 5.0.0.0          10.1.12.2                                 0 2300 500 i
```

Task 2

Using the community attribute configure R1 such that when it advertises network 1.0.0.0 /8 to R2 in AS 200, the network is not advertised to any of R2's IBGP or EBGP neighbors.

The community attribute is a numerical value that can be attached to a given prefix and advertised to a specific neighbor, once the neighbor receives the prefix, it will examine the community value and it will perform either filtering or use that value for route selection process.

By default no community attribute is sent to any neighbor. To specify that a community attribute should be sent to a BGP neighbor, the "neighbor send-community" command is configured in the router config mode.
The well known communities are as follows:

- **Internet – If assigned to a network/s, that network/s should be advertised.**
- **Local-as – If assigned to a network/s, that network/s should ONLY be advertised within that AS.**
- **No-advertise – If assigned to a network/s, that network/s should NOT be advertised to any BGP neighbor.**
- **No-export – If assigned to a network/s, that network/s should NOT be advertised to an EBGP neighbor.**

On R1

Note before configuring an access-list, always perform a "Show access-list" command to ensure that an existing access-list wil not get overridden.

R1(config)#access-list 1 permit 1.0.0.0 0.255.255.255

The access-list is used to identify the network; Prefix-lists can also be used fro this purpose.

```
R1(config)#route-map TEST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set community no-advertise

R1(config-route-map)#route-map TEST permit 20
```

Note the above route-map matches on the access-list and sets the community to one of the well known community attribute of "no-advertise", this well known community attribute tells the receiving router NOT to advertise the prefix to any of it's neighbor/s (IBGP or EBGP).
The "route-map TST permit 20" is the catch-all route-map; it basically matches any network not matched with the match keyword in the "route-map TST permit 10".

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 send-community
R1(config-router)#neighbor 10.1.12.2 route-map TEST out
```

In the above commands, we are sending the community and assigning the route-map to a given neighbor in the out direction.

The direction of the route-map specifies which routers decision should be influenced by this policy, if it should affect neighbor's decision, then, the direction of the route-map should be "out", but if the local router's decision should be influenced, then, the direction of the route-map should be "in".

To verify the configuration:

On R2

R2#Sh ip bgp 1.0.0.0

```
BGP routing table entry for 1.0.0.0/8, version 8
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to any peer)
  Not advertised to any peer
  100
    10.1.12.1 from 10.1.12.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-advertise
```

Note the community attribute from R2's perspective. Since R2 does NOT advertise the network R3 and the other EBGP neighbor won't have any knowledge of this route.

On R3

R3#Sh ip bgp 1.0.0.0

% Network not in table

On R3

R3#Show ip bgp

```
BGP table version is 8, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.24.4                0    100      0 400 i
*> 5.0.0.0          10.1.35.5                0             0 500 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 7, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.24.2                0             0 2300 i
*> 3.0.0.0          10.1.24.2                              0 2300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.24.2                              0 2300 500 i
```

Task 3

Configure R5 such that when it advertises its network 5.0.0.0 to R3 in AS 2300, the routers in AS 2300 do NOT advertise that network to any of their EBGP peer/s. DO NOT configure R3 to accomplish this task.

On R5

Note before configuring an access-list, always perform a "Show access-list" command to ensure that an existing access-list wil not get overridden.

R5(config-router)#access-list 1 permit 5.0.0.0

R5(config)#route-map TST permit 10
R5(config-route-map)#match ip addr 1
R5(config-route-map)#set community **no-export**
R5(config-route-map)# route-map TST permit 20

R5(config)#router bgp 500
R5(config-router)#neighbor 10.1.35.3 send-community
R5(config-router)#neighbor 10.1.35.3 route-map TST out

This is another well known community. In this case network 5.0.0.0 will ONLY be advertised to the routers in AS 2300. The routers in AS 2300 will NOT advertise this network to any of their EBGP neighbor/s. BUT REMEMBER THAT BY DEFAULT ROUTERS WILL STRIP THE COMMUNITY ATTRIBUTE, therefore, in this case R3 should be configured to send community to R2, or else R2 will advertise that network to its EBGP peers.

To verify the configuration:

On R3

R3(config)#Router bgp 2300
R3(config-router)#**Neighbor 10.1.23.2 send-community**

R3#Sh ip bgp 5.0.0.0

```
BGP routing table entry for 5.0.0.0/8, version 8
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to EBGP peer
Flag: 0x880
  Advertised to update-groups:
    1
  500
    10.1.35.5 from 10.1.35.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-export
```

To test this configuration further, a point-to-point frame-relay connection and an EBGP peer session can be established between R3 and R4. R4 should NOT receive an update for network 5.0.0.0 from R3, but R4 will receive an update for network 5.0.0.0 from R2.

Task 4

Configure R3 in AS 2300 to advertise network 3.0.0.0 /8 to the routers in its own AS ONLY, R3 should NOT advertise this network to any of its EBGP peers.

On R3

Note before configuring an access-list, always perform a "Show access-list" command to ensure that an existing access-list wil not get overridden.

```
R3(config)#route-map TST permit 10
R3(config-route-map)#set community local-as

R3(config)#Router bgp 2300
R3(config-router)#Network 3.0.0.0 route-map TST
R3(config-router)#Neighbor 10.1.23.2 send-community
```

Note in this case the "route-map TEST" command is applied to the inbound, because it should affect the local router for that network and community.

To verify the configuration:

On R3

R3#Sh ip bgp 3.0.0.0

```
BGP routing table entry for 3.0.0.0/8, version 5
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised outside local
 AS)
Flag: 0x820
  Advertised to update-groups:
     1
  Local
    0.0.0.0 from 0.0.0.0 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
      Community: local-AS
```

R5#Sh ip bgp

```
BGP table version is 16, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.35.3                              0 2300 i
*> 4.0.0.0          10.1.35.3                              0 2300 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

R1 is advertising network 1.0.0.0 which has an attached community attribute of "no-advertise" to R2 (Task 2). Router R2 should be configured to advertise network 1.0.0.0 to all of its IBGP and EBGP peers. You should utilize a well known community attribute to accomplish this task.

On R2

Note before configuring R2, we should display the prefix in BGP as follows:

R2#Show ip bgp 1.0.0.0

```
BGP routing table entry for 1.0.0.0/8, version 6
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to any peer)
  Not advertised to any peer
  100
    10.1.12.1 from 10.1.12.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-advertise
```

R2 can be configured to assign a well known community of "Internet" to this network, when the "Internet" community is assigned to a network, that network will be advertised to all peers.

To configure:

Note before configuring an access-list, always perform a "Show access-list" command to ensure that an existing access-list wil not get overridden.

R2(config)#access-list 1 permit 1.0.0.0

R2(config)#route-map TST permit 10

R2(config-route-map)#match ip addr 1
R2(config-route-map)#set community **Internet**

R2(config)#router bgp 2300
R2(config-router)#neighbor 10.1.12.1 route-map TST in

To verify the configuration:

On R2

R2#Show ip bgp 1.0.0.0

```
BGP routing table entry for 1.0.0.0/8, version 2
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1            2
  100
    10.1.12.1 from 10.1.12.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: internet
```

On R4

R4#Show ip bgp

```
BGP table version is 22, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.24.2                              0 2300 100 i
*> 2.0.0.0          10.1.24.2                0             0 2300 i
*> 3.0.0.0          10.1.24.2                              0 2300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.24.2                              0 2300 500 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 14, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.24.4                0    100      0 400 i
*> 5.0.0.0          10.1.35.5                0             0 500 i
```

Task 6

Erase the startup config and reload the routers before proceeding to the next task.

Task 7

Configure the routers according to the diagram/chart below and use the IP addressing and AS numbering identified in the chart.

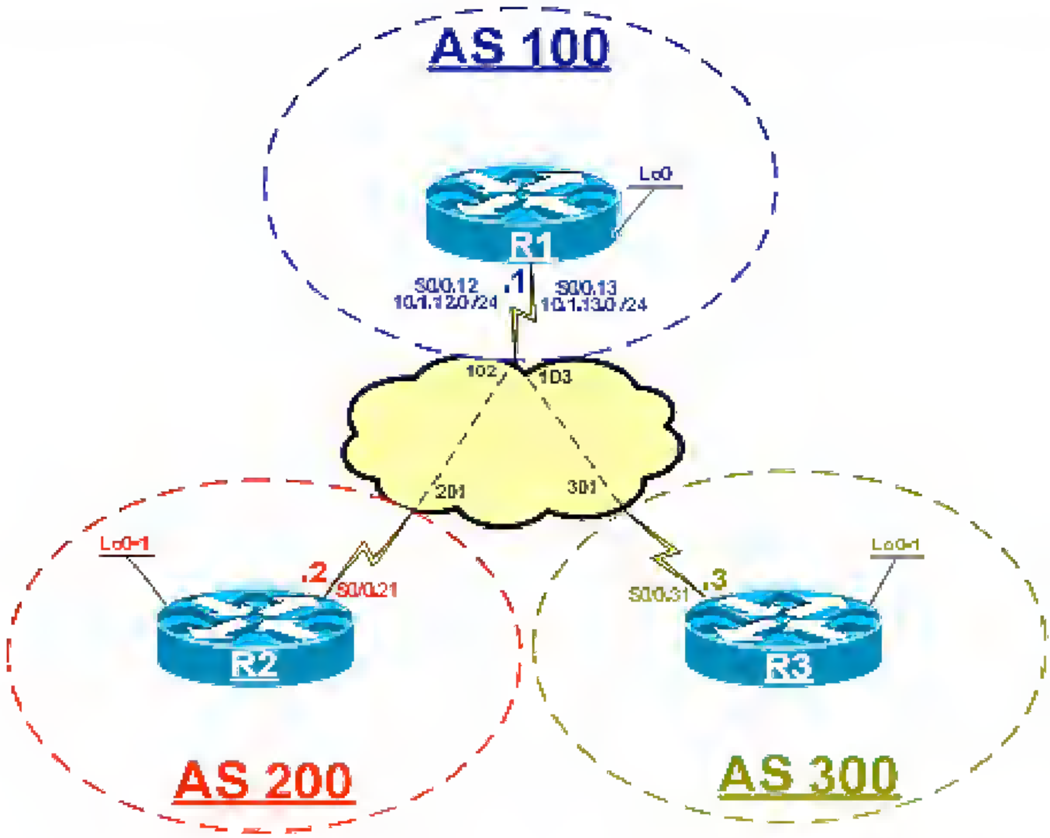

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 should have two point-to-point sub-interfaces, one connecting to R2 and the other connecting to R3.
- R2 and R3 should be configured with a single point-to-point connection to R1

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12 S0/0.13	1.1.1.1 /8 10.1.12.1 /24 10.1.13.1 /24	100
R2	**Lo0** **Lo1** **S0/0.21**	**20.1.2.2 /24** **20.1.3.2 /24** **10.1.12.2 /24**	**200**
R3	Lo0 Lo1 S0/0.31	30.1.2.3 /24 30.1.3.3 /24 10.1.13.3 /24	300

On R1

```
R1(config)#router bgp 100
R1(config-router)#netw 1.0.0.0
R1(config-router)#no au
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.13.3 remote-as 300
```

On R2

```
R2(config-if)#router bgp 200
R2(config-router)#no au
R2(config-router)#network 20.1.2.0 mask 255.255.255.0
R2(config-router)#network 20.1.3.0 mask 255.255.255.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

On R3

```
R3(config-if)#router bgp 300
R3(config-router)#no au
R3(config-router)#netw 30.1.2.0 mask 255.255.255.0
R3(config-router)#netw 30.1.3.0 mask 255.255.255.0
3(config-router)#neighbor 10.1.13.1 remote-as 100
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0              0         32768 i
*> 20.1.2.0/24      10.1.12.2            0             0 200 i
*> 20.1.3.0/24      10.1.12.2            0             0 200 i
*> 30.1.2.0/24      10.1.13.3            0             0 300 i
*> 30.1.3.0/24      10.1.13.3            0             0 300i
```

Task 8

Ensure that R1 uses AS 200 to connect to networks in subnet 2 (20.1.**2**.0 /24 and 30.1.**2**.0 /24) and AS 300 to connect to networks in subnet 3 (20.1.**3**.0 /24 and 30.1.**3**.0 /24). You must use community tags in AS 200 and 300 and neighbor commands on R1 to accomplish this task.

On R2

```
R2(config)#access-list 2 permit 20.1.2.0 0.0.0.255

R2(config)#access-list 3 permit 20.1.3.0 0.0.0.255

R2(config)#route-map TST permit 10
R2(config-route-map)#match ip addr 2
R2(config-route-map)#set community 2

R2(config)#route-map TST permit 20
R2(config-route-map)#match ip addr 3
R2(config-route-map)#set community 3
```

Note the above command "set community" tags the route/s identified in the access-list.

```
R2(config)#route-map TST permit 30

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 send-community
```

R2(config-router)#**neighbor 10.1.12.1 route-map TST out**

On R3

R3(config)#access-list 2 permit 30.1.2.0 0.0.0.255

R3(config)#access-list 3 permit 30.1.3.0 0.0.0.255

R3(config)#route-map TST permit 10
R3(config-route-map)#match ip addr 2
R3(config-route-map)#set community 2

R3(config)#route-map TST permit 20
R3(config-route-map)#match ip addr 3
R3(config-route-map)#set community 3

R3(config)#route-map TST permit 30

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.13.1 send-community
R3(config-router)#neighbor 10.1.13.1 route-map TST out

Note R1 can display the routes via their assigned community tags:

On R1

R1#Show ip bgp community 2

```
BGP table version is 18, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.1.2.0/24      10.1.12.2                0             0 200 i
*> 30.1.2.0/24      10.1.13.3                0             0 300 i
```

Note these are the routes that R2 and R3 tagged using community 2.

R1#Show ip bgp community 3

```
BGP table version is 18, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 20.1.3.0/24      10.1.12.2                0             0 200 i
*> 30.1.3.0/24      10.1.13.3                0             0 300 i
```

Note these are the routes that R2 and R3 tagged using community 3.

On R1

R1(config)#ip community-list standard TST2 permit 2

R1(config)#ip community-list standard TST3 permit 3

R1 is identifying the community tags using a community-list. This is like writing an access-list to identify a given route/s.

R1(config)#route-map TEST permit 10
R1(config-route-map)#match community TST2
R1(config-route-map)#set ip next-hop 10.1.12.2

R1(config)#route-map TEST permit 20
R1(config-route-map)#match community TST3
R1(config-route-map)#set ip next-hop 10.1.13.3

R1(config)#route-map TEST permit 30

The communities are matched and the policy is assigned.

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-map TEST in
R1(config-router)#neighbor 10.1.13.3 route-map TEST in

The policies are applied to the neighbors using the "neighbor route-map" commands.
Do not try to test reachability to the network through the newly assigned next hop IP addresses, the purpose of this lab is to understand the scope of the community attributes and its uses. To verify this lab enter "Show ip bgp" to see the next hop attribute.

R1#Show ip bgp

BGP table version is 6, local router ID is 1.1.1.1

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf      Weight   Path
*> 1.0.0.0           0.0.0.0            0                32768   i
*> 20.1.2.0/24      10.1.12.2           0                    0   200 i
*> 20.1.3.0/24      10.1.13.3           0                    0   200 i
*> 30.1.2.0/24      10.1.12.2           0                    0   300 i
*> 30.1.3.0/24      10.1.13.3           0                    0   300 i
```

Task 9

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 7 – BGP Cost Community

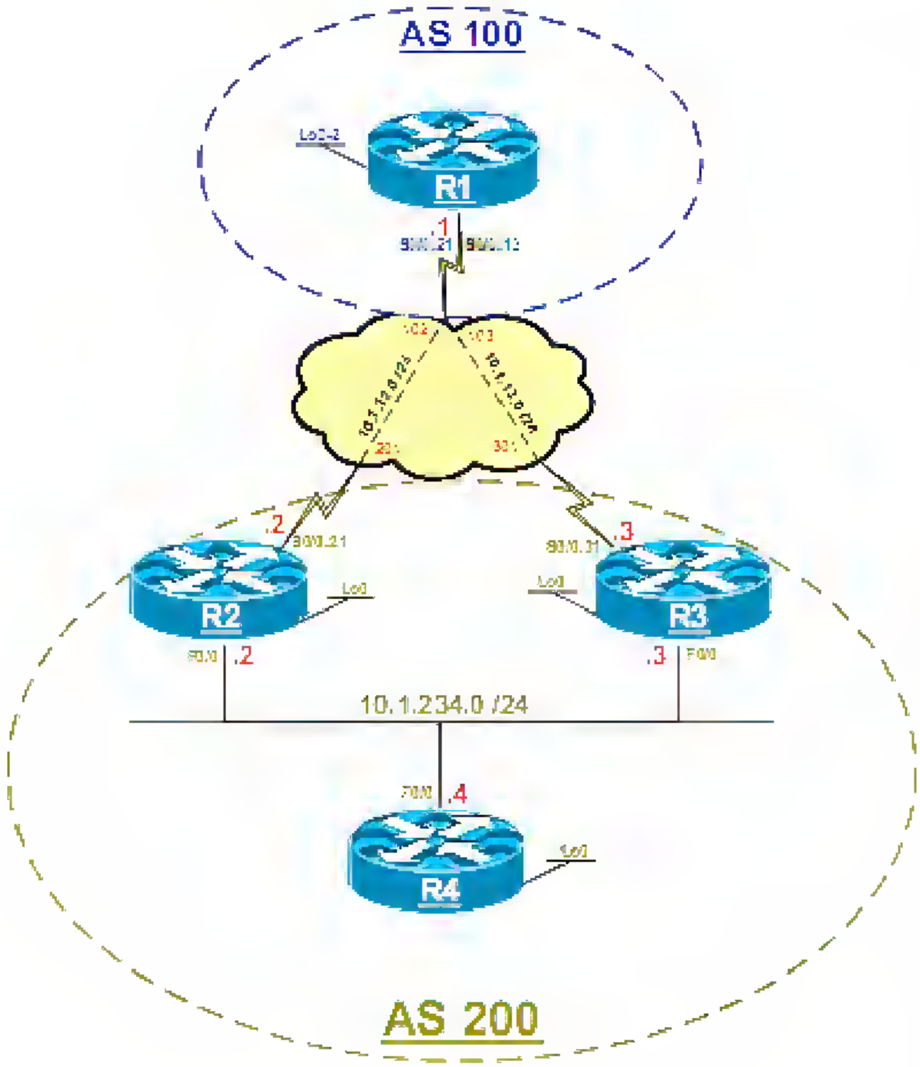

Lab Setup:

- Configure R1 to be the hub and R2 and R3 to be the spokes, the frame-relay routers should be configured in a point-to-point manner.
- RIPv2 should be used to provide NLRI for the links. The loopback0 interfaces of R2, R3 and R4 should also be advertised in RIPv2 routing protocol.
- The F0/0 interface of R2, R3 and R4 should be configured in VLAN 234.

➢ Use the following IP addressing chart for IP addressing assignment

IP addressing:

Router	Interface / IP address
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24 Loopback 0 = 1.1.1.1 /8 Loopback 1 = 100.1.1.1 /24 Loopback 2 = 200.1.1.1 /24
R2	**S0/0.21 = 10.1.12.2 /24** **F0/0 = 10.1.234.2 /24** **Loopback0 = 2.2.2.2 /8**
R3	S0/0.31 = 10.1.13.3 /24 F0/0 = 10.1.234.3 /24 Loopback0 = 3.3.3.3 /8
R4	**F0/0 = 10.1.234.4 /24** **Loopback0 = 4.4.4.4 /8**

Task 1

Configure R1 in AS 100, this router should establish an EBGP peer session with R2 and R3 in AS 200, R1 should advertise it's Lo1 and Lo2 interface in BGP. All BGP routers should use their loopback 0 interface as their router-id.

On R1

```
R1(config)#router bgp 100
R1(config-router)#bgp router-id 1.1.1.1
R1(config-router)#no auto-summary

R1(config-router)#neigh 10.1.12.2 remote-as 200
R1(config-router)#neigh 10.1.13.3 remote-as 200
```

Note in BGP if the auto summary is disabled, then a subnetted network should be advertised using the mask keyword followed by the correct mask.

R1(config-router)#network 100.1.1.0 mask 255.255.255.0
R1(config-router)#network 200.1.1.0

On R2

R2(config)#router bgp 200
R2(config-router)#bgp router-id 2.2.2.2
R2(config-router)#no auto-summary

R2(config-router)#neighbor 10.1.12.1 remote-as 100

On R3

R3(config)#router bgp 200
R3(config-router)#bgp router-id 3.3.3.3
R3(config-router)#no auto-summary

R3(config-router)#neighbor 10.1.13.1 remote-as 200

To verify the configuration:

On R1

R1#Show ip bgp summary

BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 3, main routing table version 3
2 network entries using 234 bytes of memory
2 path entries using 104 bytes of memory
2/1 BGP path/bestpath attribute entries using 248 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 586 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.12.2	4	200	4	5	3	0	0	00:01:34	0
10.1.13.3	4	200	4	5	3	0	0	00:00:20	0

R1#Show ip bgp

BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

```
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop            Metric  LocPrf Weight  Path
*> 100.1.1.0/24     0.0.0.0                  0         32768  i
*> 200.1.1.0        0.0.0.0                  0         32768  i
```

On R2

R2#Show ip bgp

```
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop            Metric  LocPrf Weight  Path
*> 100.1.1.0/24     10.1.12.1                0             0  100 i
*> 200.1.1.0        10.1.12.1                0             0  100 i
```

On R3

R3#Show ip bgp

```
BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop            Metric  LocPrf Weight  Path
*> 100.1.1.0/24     10.1.13.1                0             0   100 i
*> 200.1.1.0        10.1.13.1                0             0   100 i
```

Task 2

Configure an IBGP peer session between R2, R3 and R4; these routers should establish their peer session based on their Loopback 0 interface.

On R2

```
R2(config)#router bgp 200
R2(config-router)#neighbor 3.3.3.3 remote-as 200
```

R2(config-router)#neighbor 3.3.3.3 update-source lo0
R2(config-router)#neighbor 4.4.4.4 remote-as 200
R2(config-router)#neighbor 4.4.4.4 update-source lo0

On R3

R3(config)#router bgp 200
R3(config-router)#neighbor 2.2.2.2 remote-as 200
R3(config-router)#neighbor 2.2.2.2 update-source lo0

R3(config-router)#neighbor 4.4.4.4 remote-as 200
R3(config-router)#neighbor 4.4.4.4 update-source lo0

On R4

R4(config)#router bgp 200
R4(config-router)#no auto-summary
R4(config-router)#bgp router-id 4.4.4.4

R4(config-router)#neighbor 2.2.2.2 remote-as 200
R4(config-router)#neighbor 2.2.2.2 update-source lo0

R4(config-router)#neighbor 3.3.3.3 remote-as 200
R4(config-router)#neighbor 3.3.3.3 update-source lo0

To verify the configuration:

On R2

R2#Show ip bgp summ

BGP router identifier 2.2.2.2, local AS number 200
BGP table version is 3, main routing table version 3
<snip>

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.3.3.3	4	200	7	7	3	0	0	00:02:44	2
4.4.4.4	4	200	5	6	3	0	0	00:01:27	0
10.1.12.1	4	100	31	30	3	0	0	00:27:38	2

On R3

R3#Show ip bgp summ

BGP router identifier 3.3.3.3, local AS number 200

```
BGP table version is 3, main routing table version 3
2 network entries using 234 bytes of memory
4 path entries using 208 bytes of memory
3/1 BGP path/bestpath attribute entries using 372 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 838 total bytes of memory
BGP activity 2/0 prefixes, 4/0 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4   200      12      12        3    0    0 00:07:55        2
4.4.4.4         4   200       9      10        3    0    0 00:05:59        0
10.1.13.1       4   100      36      35        3    0    0 00:31:36        2
```

On R4

R4#Show ip bgp summ

```
BGP router identifier 4.4.4.4, local AS number 200
BGP table version is 3, main routing table version 3
2 network entries using 234 bytes of memory
4 path entries using 208 bytes of memory
2/1 BGP path/bestpath attribute entries using 248 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 714 total bytes of memory
BGP activity 2/0 prefixes, 4/0 paths, scan interval 60 secs


Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4   200      15      14        3    0    0 00:10:40        2
3.3.3.3         4   200      15      14        3    0    0 00:10:01        2
```

Task 3

Configure R2 and R3 to result the following output:

On R2

R2#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24    3.3.3.3                  0    100      0 100 i
*>                 10.1.12.1                0             0 100 i
* i200.1.1.0       3.3.3.3                  0    100      0 100 i
*>                 10.1.12.1                0             0 100 i
```

On R3

R3#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24    2.2.2.2                  0    100      0 100 i
*>                 10.1.13.1                0             0 100 i
* i200.1.1.0       2.2.2.2                  0    100      0 100 i
*>                 10.1.13.1                0             0 100 i
```

On R4

R4#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24    3.3.3.3                  0    100      0 100 i
*>i                2.2.2.2                  0    100      0 100 i
* i200.1.1.0       3.3.3.3                  0    100      0 100 i
*>i                2.2.2.2                  0    100      0 100 i
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#neighbor 3.3.3.3 next-hop-self
R2(config-router)#neighbor 4.4.4.4 next-hop-self
```

On R3

```
R3(config)#router bgp 200
R3(config-router)#neighbor 2.2.2.2 next-hop-self
R3(config-router)#neighbor 4.4.4.4 next-hop-self
```

To verify the configuration:

On R2

R2#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24    3.3.3.3                  0    100      0 100 i
*>                 10.1.12.1                0             0 100 i
* i200.1.1.0       3.3.3.3                  0    100      0 100 i
*>                 10.1.12.1                0             0 100 i
```

On R3

R3#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24    2.2.2.2                  0    100      0 100 i
*>                 10.1.13.1                0             0 100 i
* i200.1.1.0       2.2.2.2                  0    100      0 100 i
*>                 10.1.13.1                0             0 100 i
```

On R4

R4#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24    3.3.3.3                  0    100      0 100 i
*>i                2.2.2.2                  0    100      0 100 i
* i200.1.1.0       3.3.3.3                  0    100      0 100 i
*>i                2.2.2.2                  0    100      0 100 i
```

Task 4

Configure R4 such that if a "Show ip bgp | b Network" command is entered, it matches the output of the following:

On R4

R4#Show ip bgp | b Network

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i100.1.1.0/24	3.3.3.3	0	100	0	100 i
* i	2.2.2.2	0	100	0	100 i
* i200.1.1.0	3.3.3.3	0	100	0	100 i
*>i	2.2.2.2	0	100	0	100 i

In this scenario, the cost extended community attribute is used.

- Cost is an extended community attribute

- It's a Non-Transitive extended community attribute that allows you to customize the local route preference which can influence the best path selection process

- This attribute is applied by configuring the "Set extcommunity cost" command, using a route-map. This command is configured with a cost community id (0-255) and a cost value (0-4.29 Billion) with a default cost value of 2.145 Billion. The lower value has more preference, but the lower cost community id value is used as the tie breaker.

On R4

The following identifies the prefix:

R4(config)#**access-list 1 permit 100.1.1.0 0.0.0.255**

A route-map matches the access-list and applies the extcommunity cost attribute with two numbers, the first number is the community id and the second number is the community value.

R4(config)#Route-map TST permit 10
R4(config-route-map)#**match ip address 1**
R4(config-route-map)#**Set extcommunity cost** 1 1

R4(config-route-map)#route-map TST per 20

Lastly, it's applied by the neighbor command:

R4(config-route-map)#router bgp 200
R4(config-router)#**neighbor 3.3.3.3 route-map TST in**

Task 5

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 8 – BGP & Load Balancing-I

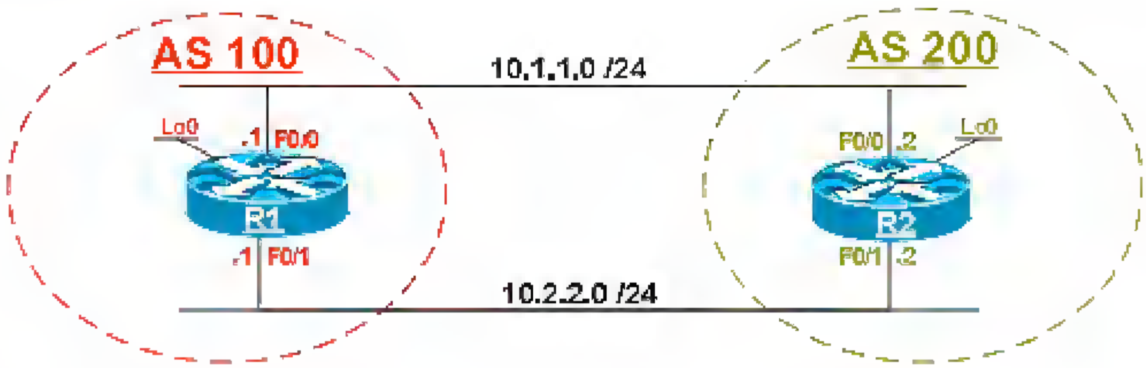

Lab Setup:

- Configure F0/0 interface of R1 and R2 in VLAN 100
- Configure F0/1 interface of R1 and R2 in VLAN 200
- Use the IP addressing chart for IP address assignment

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 F0/1 = 10.2.2.1 /24 Loopback 0 = 1.1.1.1 /8
R2	F0/0 = 10.1.1.2 /24 F0/0 = 10.2.2.2 /24 Loopback0 = 2.2.2.2 /8

Task 1

Configure an EBGP peer session between R1 and R2 ensure that these routers perform load balancing using the two links. Use an IGP of your choice.

In this topology since the routers are directly connected, the load balancing can be performed if the EBGP peer session is established based on the loopback interface of the routers, up to 6 equal cost paths can be used. RIPv2 was chosen as the IGP.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no auto-summary

R1(config-router)#neighbor 2.2.2.2 remote-as 200
R1(config-router)#neighbor 2.2.2.2 ebgp-multihop 2
R1(config-router)#neighbor 2.2.2.2 update-source lo0

R1(config)#router rip
R1(config-router)#no auto-summary
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
R1(config-router)#network 1.0.0.0
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no auto-summary

R2(config-router)#neighbor 1.1.1.1 remote-as 100
R2(config-router)#neighbor 1.1.1.1 update-source lo0
R2(config-router)#neighbor 1.1.1.1 ebgp-multihop 2

R2(config)#router rip
R2(config-router)#no auto-summary
R2(config-router)#ver 2
R2(config-router)#network 10.0.0.0
R2(config-router)#network 2.0.0.0
```

To verify the configuration:

On R1

R1#Sh ip bgp summ

```
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4   200       7       8        1    0    0 00:01:07        0
```

On R1

R1#Show ip route rip

```
R    2.0.0.0/8 [120/1] via 10.2.2.2, 00:00:03, FastEthernet0/1
               [120/1] via 10.1.1.2, 00:00:27, FastEthernet0/0
```

On R2

R2#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.2.2.1, 00:00:19, FastEthernet0/1
               [120/1] via 10.1.1.1, 00:00:03, FastEthernet0/0
```

To test the configuration:

On R1

R1#Traceroute 2.2.2.2

```
Type escape sequence to abort.
Tracing the route to 2.2.2.2

  1 10.2.2.2 4 msec
    10.1.1.2 4 msec *
```

On R2

R2#Traceroute 1.1.1.1

```
Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.2.2.1 0 msec
    10.1.1.1 0 msec *
```

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 9 – BGP & Load Balancing-II

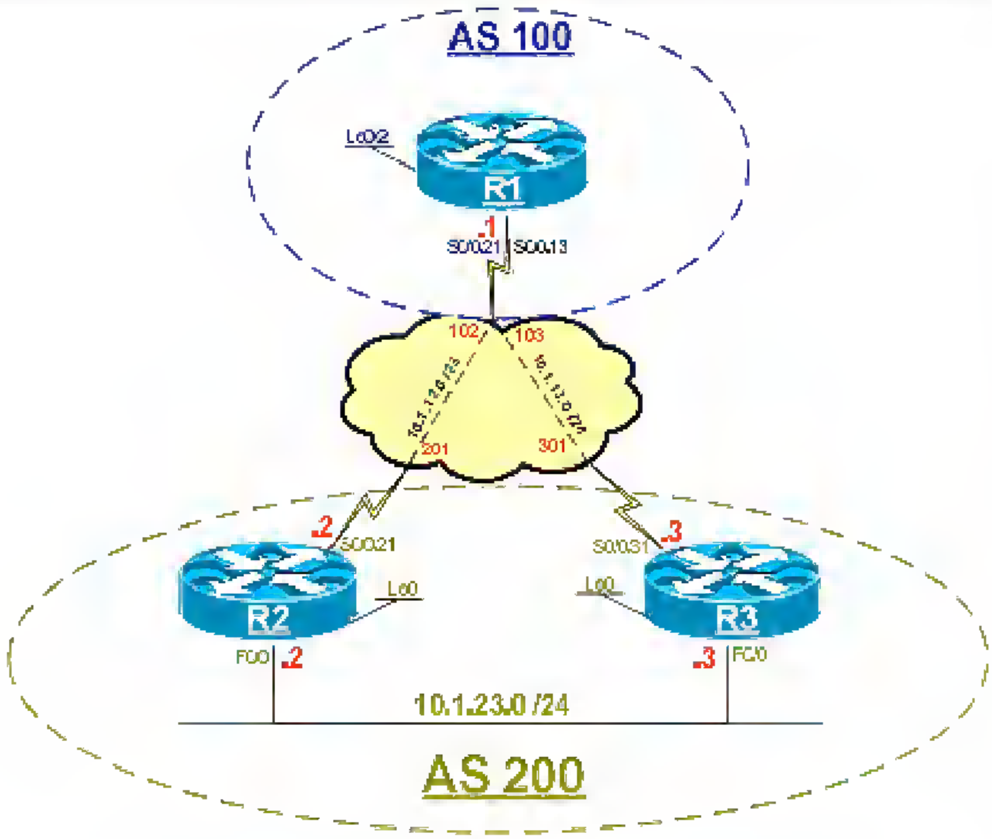

Lab Setup:

- Configure R1 as the hub and R2 and R3 as the spokes, all frame-relay links should be configured in a point-to-point manner.
- Configure F0/0 interface of R2 and R3 in VLAN 23.
- Use the IP addressing chart for IP addressing assignment.

IP addressing:

Router	Interface / IP address
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24 Loopback 0 = 1.1.1.1 /8
R2	S0/0.21 = 10.1.12.2 /24 F0/0 = 10.1.23.2 /24 Loopback0 = 2.2.2.2 /8
R3	S0/0.31 = 10.1.13.3 /24 F0/0 = 10.1.23.3 /24 Loopback0 = 3.3.3.3 /8

Task 1

Configure R1 in AS 100 to establish EBGP peer sessions with R2 and R3 in AS 300. R2 and R3 should advertise netw 10.1.23.0 /24 in BGP.

On R1

R1(config-router)#router bgp 100
R1(config-router)#no auto-summary

R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.13.3 remote-as 200

On R2

R2(config)#router bgp 200
R2(config-router)#no auto-summary
R2(config-router)#Network 10.1.23.0 mask 255.255.255.0

R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200

On R3

R3(config)#router bgp 200
R3(config-router)#no auto-summary
R3(config-router)#network 10.1.23.0 mask 255.255.255.0
R3(config-router)#neighbor 10.1.13.1 remote-as 100

```
R3(config-router)#neighbor 10.1.23.2 remote-as 200
```

To verify the configuration:

On R1

```
R1#Sh ip bgp

BGP table version is 2, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  10.1.23.0/24     10.1.12.2                0             0 200 i
*>                  10.1.13.3                0             0 200 i
```

```
R1#Show ip route | b Gateway

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 3 subnets
C       10.1.13.0 is directly connected, Serial0/0.13
C       10.1.12.0 is directly connected, Serial0/0.12
B       10.1.23.0 [20/0] via 10.1.13.3, 00:03:32
```

Task 2

Configure R1 such that it uses both neighbors (R2 and R3) to perform an equal cost load balancing.

Note BGP will ONLY use one path to a given destination; therefore, it does not perform load balancing amongst multiple equal cost paths. The "maximum-paths" command can be configured to change this behavior.

On R1

```
R1(config)#router bgp 100
R1(config-router)#maximum-paths 2
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop         Metric LocPrf Weight Path
*> 10.1.23.0/24     10.1.12.2             0             0 200 i
*                   10.1.13.3             0             0 200 i
```

R1#Show ip route | b Gateway

```
Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 3 subnets
C       10.1.13.0 is directly connected, Serial0/0.13
C       10.1.12.0 is directly connected, Serial0/0.12
B       10.1.23.0 [20/0] via 10.1.13.3, 00:00:18
                  [20/0] via 10.1.12.2, 00:00:18
```

Note R1 is performing equal cost load balancing across the two links.

Task 3

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 10 – BGP Unequal-Cost Load Balancing

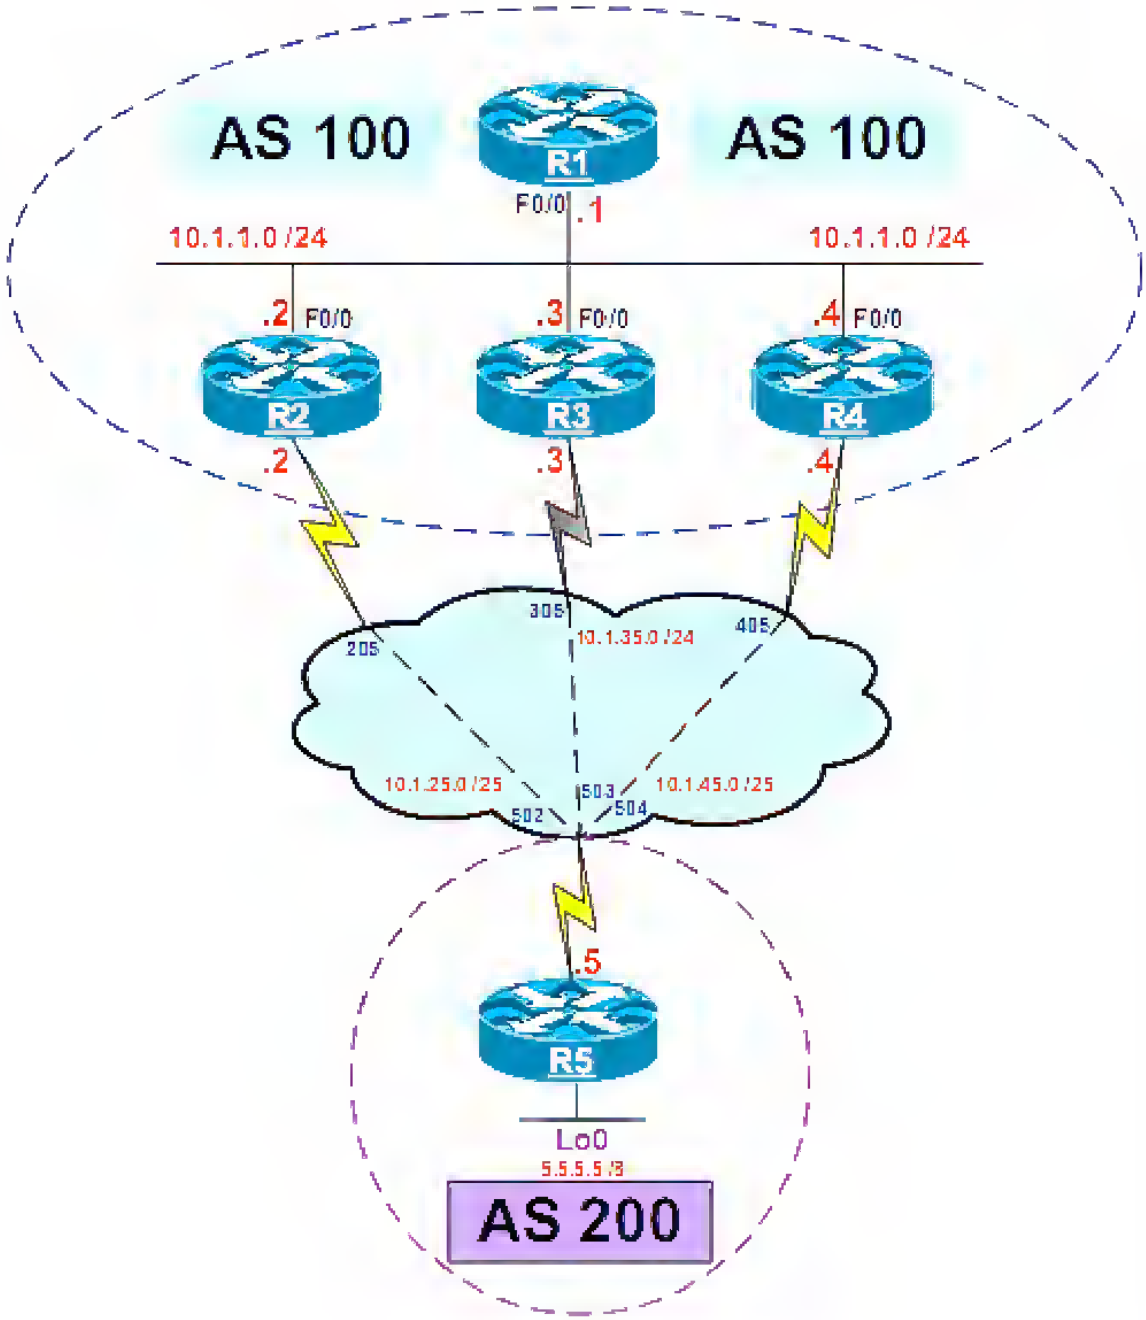

Lab Setup:

- Configure F0/0 interface of R1, R2, R3 and R4 should be configured in VLAN 100.
- Configure the frame-relay connections in a point-to-point manner.
- Use the IP addressing chart below for IP address assignment.

IP addressing:

Router	Interface / IP address	VLAN
R1	F0/0 = 10.1.1.1 /24	100
R2	**F0/0 = 10.1.1.2 /24** **S0/0.25 = 10.1.25.2 /24**	**100**
R3	S0/0.35 = 10.1.35.3 /24 F0/0 = 10.1.1.3 /24	 100
R4	**F0/0 = 10.1.1.4 /24** **S0/0.45 = 10.1.45.4 /24**	**100**
R5	S0/0.54 = 10.1.45.5 /24 S0/0.53 = 10.1.35.5 /24 S0/0.52 = 10.1.25.5 /24 Lo0 = 5.5.5.5 /8	

Task 1

Configure peering according to the diagram.

On R1

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote 100
R1(config-router)#neighbor 10.1.1.3 remote 100
R1(config-router)#neighbor 10.1.1.4 remote 100
```

On R2

```
R2(config)#router bgp 100
```

```
R2(config-router)#no au
R2(config-router)#neighbor 10.1.1.1 remote 100
R2(config-router)#neighbor 10.1.1.3 remote 100
R2(config-router)#neighbor 10.1.1.4 remote 100
R2(config-router)#neighbor 10.1.25.5 remote 200
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#no au
R3(config-router)#neighbor 10.1.1.1 remote 100
R3(config-router)#neighbor 10.1.1.2 remote 100
R3(config-router)#neighbor 10.1.1.4 remote 100
R3(config-router)#neighbor 10.1.35.5 remote 200
R3(config-router)#neighbor 10.1.35.5 remote 200
```

On R4

```
R4(config)#router bgp 100
R4(config-router)#no au
R4(config-router)#neighbor 10.1.1.1 remote 100
R4(config-router)#neighbor 10.1.1.2 remote 100
R4(config-router)#neighbor 10.1.1.3 remote 100
R4(config-router)#neighbor 10.1.45.5 remote 200
```

On R5

```
R5(config)#router bgp 200
R5(config-router)#no au
R5(config-router)#neighbor 10.1.25.2 remote 100
R5(config-router)#neighbor 10.1.35.3 remote 100
R5(config-router)#neighbor 10.1.45.4 remote 100
R5(config-router)#network 5.0.0.0
```

To verify the configuration:

On R1

R1#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0          10.1.25.5                0    100      0 200 i
* i                 10.1.35.5                0    100      0 200 i
```

Task 2

Configure the border routers to change the next hop IP address to an internal IP address.

On R2, R3 and R4

R2(config)#router bgp 100
R2(config-router)#**neighbor 10.1.1.1 next-hop-self**

To verify the configuration:

On R1

R1#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0          10.1.1.2                 0    100      0 200 i
* i                 10.1.1.3                 0    100      0 200 i
```

Task 3

Configure the routers in AS 100 such that R1 distributes traffic proportionally over the external links to reach prefix 5.0.0.0 /8, the load balancing should be done based on the bandwidth of the links between the border routers of this AS and AS 200.

The unequal cost load balancing feature is used in conjunction with BGP multipath feature to advertise the exit link's bandwidth as an extended community to IBGP peers, this feature is used for links between directly connected EBGP neighbors and available in IOS release 12.2(2T) or better.

To configure this feature, the following steps should be performed:

1. **Enable the BGP dmzlink-bw feature:**
   This is accomplished by configuring the "BGP dmzlink-bw" router configuration command, this must be configured on the border routers and the internal routers

2. **Configure BGP to include the link bandwidth value of the external interface in extended community so they can be propagated to IBGP peers. This is accomplished through the "Neighbor dmzlink-bw" router configuration mode command.**

Remember, for this feature to work, R1 must have an equal IGP cost and BGP attributes or else the feature will NOT work. Note the bandwidth of the links connecting the border routers R2, R3 and R4 to AS 200 is set to 64K, 128K and 192K respectively.

On R1

R1#Show ip bgp 5.0.0.0

```
BGP routing table entry for 5.0.0.0/8, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
Flag: 0x820
  Not advertised to any peer
  200
    10.1.1.3 from 10.1.1.3 (10.1.35.3)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    10.1.1.2 from 10.1.1.2 (10.1.25.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

R1#Show ip route 5.0.0.0

```
Routing entry for 5.0.0.0/8
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Last update from 10.1.1.2 00:00:30 ago
  Routing Descriptor Blocks:
  * 10.1.1.2, from 10.1.1.2, 00:00:30 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
```

Note BGP table identifies ONLY one of the routes as the best, in this case since all the attributes are equal, and the "BGP Bestpath compare-routerid" command is NOT configured, the neighbor with the lowest IP address was chosen.

To configure the task:

On R1

The following allows the local router to have three equal IBGP cost paths:

```
R1(config)#Router bgp 100
R1(config-router)#maximum-path ibgp 3
```

On R2, R3 and R4

Rx(config)#router bgp 100
Rx(config-router)#**neighbor 10.1.1.1 send-community extended**
Rx(config-router)#**bgp dmzlink-bw**

On R2

R2(config)#router bgp 100
R2(config-router)#**neighbor 10.1.25.5 dmzlink-bw**

On R3

R3(config)#router bgp 100
R3(config-router)#**neighbor 10.1.35.5 dmzlink-bw**

On R4

R4(config)#router bgp 100
R4(config-router)#**neighbor 10.1.45.5 dmzlink-bw**

To verify the configuration:

On R1

R1#Show ip bgp 5.0.0.0

```
BGP routing table entry for 5.0.0.0/8, version 2
Paths: (3 available, best #2, table Default-IP-Routing-Table)
Multipath: iBGP
Flag: 0x820
  Not advertised to any peer
  200
    10.1.1.2 from 10.1.1.2 (10.1.25.2)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      DMZ-Link Bw 8 kbytes
  200
    10.1.1.4 from 10.1.1.4 (10.1.1.4)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      DMZ-Link Bw 24 kbytes

  200
    10.1.1.3 from 10.1.1.3 (10.1.35.3)
```

Origin IGP, metric 0, localpref 100, valid, internal, multipath
DMZ-Link Bw 16 kbytes

R1#Show ip route 5.0.0.0

```
Routing entry for 5.0.0.0/8
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Last update from 10.1.1.3 00:04:46 ago
  Routing Descriptor Blocks:
    10.1.1.4, from 10.1.1.4, 00:04:46 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
    10.1.1.3, from 10.1.1.3, 00:04:46 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
  * 10.1.1.2, from 10.1.1.2, 00:04:46 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 200
```

Note the traffic count is 1, in order to have the BGP table reflect on this counter, the "BGP dmzlink-bw" must be configured, as follows:

On R1

R1(config)#router bgp 100
R1(config-router)#**bgp dmzlink-bw**

To verify the configuration:

On R1

R1#Show ip route 5.0.0.0

```
Routing entry for 5.0.0.0/8
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Last update from 10.1.1.3 00:00:40 ago

  Routing Descriptor Blocks:
    10.1.1.4, from 10.1.1.4, 00:00:40 ago
```

```
    Route metric is 0, traffic share count is 2
    AS Hops 1
    Route tag 200
  10.1.1.3, from 10.1.1.3, 00:00:40 ago
    Route metric is 0, traffic share count is 1

    AS Hops 1
    Route tag 200
 * 10.1.1.2, from 10.1.1.2, 00:07:53 ago
    Route metric is 0, traffic share count is 24
    AS Hops 1
    Route tag 200
```

Note the feature does not work unless it is enabled.

Task 4

Erase the startup configuration on all routers before proceeding to the next lab.

Lab 11 - Local-Preference -I

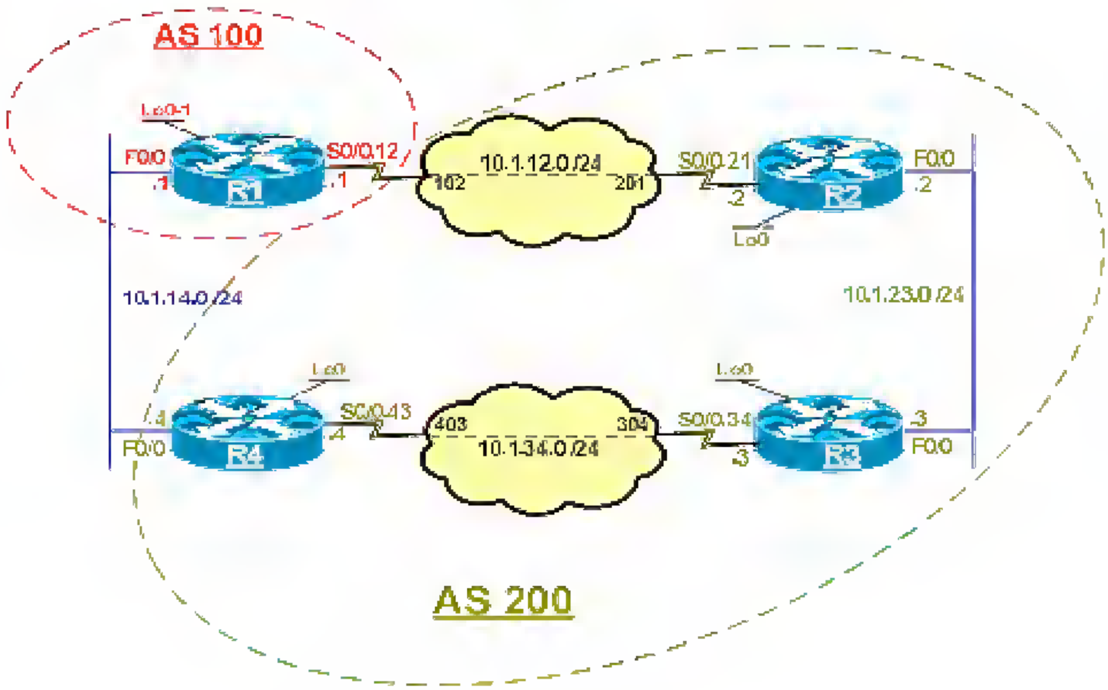

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

Ip addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	200
R4	Lo0 F0/0 S0/0.43	4.4.4.4 /8 10.1.14.4 /24 10.1.34.4 /24	200

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. Routers R2 and router R4 should have EBGP peer session to R1 in AS 100. BGP routers should ONLY advertise their loopback interface/s in BGP. Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1

R1(config-if)#router bgp 100
R1(config-router)#no au
R1(config-router)#neighbor 10.1.14.4 remote-as 200
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0

On R2

R2(config)#router bgp 200
R2(config-router)#no au

```
R2(config-router)#no syn
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
R2(config-router)#neighbor 10.1.34.4 remote-as 200

R2(config)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3

```
R3(config)#router bgp 200
R3(config-router)#no au
R3(config-router)#no syn
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 200
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4

```
R4(config)#router bgp 200
R4(config-router)#no syn
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 200
R4(config-router)#neighbor 10.1.23.2 remote-as 200
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#no au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

Task 2

Ensure that the routers in AS 200 use R4 to reach network 1.0.0.0 /8 in AS 100. Local-Pref attribute must be used to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined, as followed:

On R1

R1#Sh ip bgp

```
BGP table version is 6, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                              0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 6, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 6, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 6, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.12.1                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         10.1.12.1                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
```

Note Routers R2 and R3 are taking the R2-R1 link to connect to network 1.0.0.0 /8.

On R4

R4(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R4(config)#route-map TST permit 10
R4(config-route-map)#match ip addr 1

R4(config-route-map)#set Local-preference 400

R4(config)#route-map TST permit 20

R4(config-route-map)#router bgp 200
R4(config-router)#neighbor 10.1.14.1 route-map TST in

The local preference attribute is used to prefer an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the local AS, the local preference attribute is used to select the exit point for a specific or all routes. Since the local preference attribute affects the routers within the AS, the route-map should be configured in the "in" direction. Remember that with local preference the higher value has better preference.

To verify the configuration:

On R2

R2#Sh ip bgp

```
BGP table version is 9, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    400      0 100 i
*                   10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 10, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    400      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 8, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0    400      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 11.0.0.0         10.1.14.1                0             0 100 i
* i                 10.1.12.1                0    100      0 100 i
```

Note the routers in AS 200 take the R4-R1 link to connect to network 1.0.0.0/8.

Task 3

Ensure that the routers in AS 200 use R2 to reach network 11.0.0.0 /8. Local-Pref attribute must be used to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined, as followed:

On R2

R2#Sh ip bgp

```
BGP table version is 9, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    400      0 100 i
*                   10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 10, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    400      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 8, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0    400      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 11.0.0.0         10.1.14.1                0             0 100 i
```

```
* i                 10.1.12.1          0      100         0 100 i
```

Note router R4 is taking the R4-R1 link to connect to network 11.0.0.0.

On R2

```
R2(config)#access-list 11 permit 11.0.0.0 0.255.255.255

R2(config)#route-map TST permit 10
R2(config-route-map)#match ip addr 11
R2(config-route-map)#set Local-preference 400

R2(config)#route-map TST permit 20

R2(config-route-map)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST in
```

To verify the configuration:

On R2

R2#Sh ip bgp

```
BGP table version is 6, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          10.1.12.1                0             0 100 i
*>i                 10.1.14.1                0    400      0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*> 11.0.0.0         10.1.12.1                0    400      0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 14, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
    Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    400      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*>i11.0.0.0         10.1.12.1                0    400      0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 11, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0    400      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i11.0.0.0         10.1.12.1                0    400      0 100 i
*                   10.1.14.1                0             0 100 i
```

Note the routers in AS 200 connect to network 11.0.0.0 /8 by going through R2-R1 link.

Task 4

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 12 – BGP Local-Preference - II

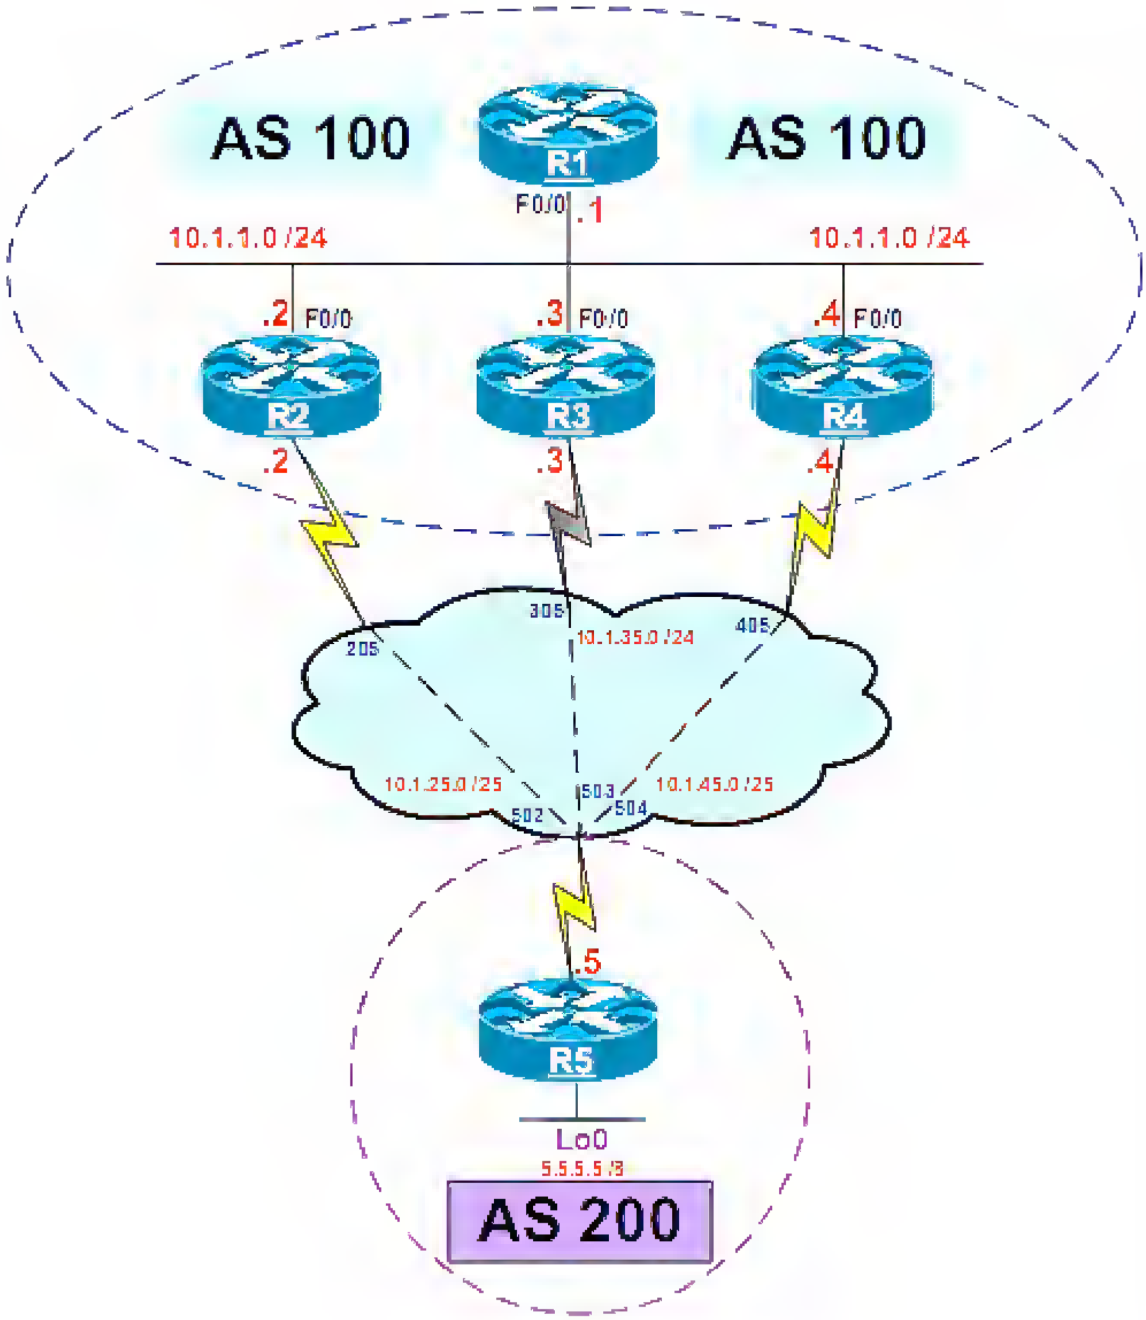

Lab Setup:

- Configure F0/0 interface of R1, R2, R3 and R4 should be configured in VLAN 100.
- Configure the frame-relay connections in a point-to-point manner.
- Use the IP addressing chart below for IP address assignment.

IP addressing:

Router	Interface / IP address	VLAN
R1	F0/0 = 10.1.1.1 /24	100
R2	**F0/0 = 10.1.1.2 /24** **S0/0.25 = 10.1.25.2 /24**	**100**
R3	S0/0.35 = 10.1.35.3 /24 F0/0 = 10.1.1.3 /24	 100
R4	**F0/0 = 10.1.1.4 /24** **S0/0.45 = 10.1.45.4 /24**	**100**
R5	S0/0.54 = 10.1.45.5 /24 S0/0.53 = 10.1.35.5 /24 S0/0.52 = 10.1.25.5 /24 Lo0 = 5.5.5.5 /8	

Task 1

Configure peering according to the diagram.

On R1

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote 100
R1(config-router)#neighbor 10.1.1.3 remote 100
R1(config-router)#neighbor 10.1.1.4 remote 100
```

On R2

```
R2(config)#router bgp 100
R2(config-router)#no au
```

R2(config-router)#neighbor 10.1.1.1 remote 100

R2(config-router)#neighbor 10.1.1.3 remote 100
R2(config-router)#neighbor 10.1.1.4 remote 100
R2(config-router)#neighbor 10.1.25.5 remote 200

On R3

R3(config)#router bgp 100
R3(config-router)#no au
R3(config-router)#neighbor 10.1.1.1 remote 100
R3(config-router)#neighbor 10.1.1.2 remote 100
R3(config-router)#neighbor 10.1.1.4 remote 100
R3(config-router)#neighbor 10.1.35.5 remote 200
R3(config-router)#neighbor 10.1.35.5 remote 200

On R4

R4(config)#router bgp 100
R4(config-router)#no au
R4(config-router)#neighbor 10.1.1.1 remote 100
R4(config-router)#neighbor 10.1.1.2 remote 100
R4(config-router)#neighbor 10.1.1.3 remote 100
R4(config-router)#neighbor 10.1.45.5 remote 200

On R5

R5(config)#router bgp 200
R5(config-router)#no au
R5(config-router)#neighbor 10.1.25.2 remote 100
R5(config-router)#neighbor 10.1.35.3 remote 100
R5(config-router)#neighbor 10.1.45.4 remote 100
R5(config-router)#network 5.0.0.0

To verify the configuration:

On R1

R1#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0          10.1.25.5                0    100      0 200 i
* i                 10.1.35.5                0    100      0 200 i
```

Task 2

Configure the border routers to change the next to an internal IP address.

On R2, R3 and R4

R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.1.1 next-hop-self

To verify the configuration:

On R1

R1#Show ip bgp | B Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*  i5.0.0.0         10.1.1.3                 0    100      0 200 i
*>i                 10.1.1.2                 0    100      0 200 i
*  i                10.1.1.4                 0    100      0 200 i
```

Task 3

Configure R2, R3 and R4 such that R1 takes R4 as the primary and R3 as the backup, if R4 and R3 are both down, then it should take R2 to reach Network 5.0.0.0 /8. You must use local preference to accomplish this task.

On R2

R2(config)#router bgp 100
R2(config-router)#bgp default local-preference 200
R2(config-router)#do cle ip bgp * out

On R3

R3(config)#router bgp 100
R3(config-router)#bgp default local-preference 300
R3(config-router)#do clc ip bgp * out

On R4

```
R4(config-router)#bgp default local-preference 400
R4(config-router)#do clc ip bgp * out
```

To verify the configuration:

On R1

R1#Show ip bgp | b Network

```
  Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0         10.1.1.4                 0    400      0 200 i
```

R1#Show ip bgp 5.0.0.0

```
BGP routing table entry for 5.0.0.0/8, version 6
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  200
    10.1.1.4 from 10.1.1.4 (10.1.45.4)
      Origin IGP, metric 0, localpref 400, valid, internal, best
```

To test the configuration:

On R1

R1#Traceroute 5.5.5.5

```
Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 10.1.1.4 0 msec 4 msec 0 msec
  2 10.1.45.5 32 msec *  28 msec
```

Task 4

Remove the configuration from the previous task and re-configure the same task using another method. DO NOT use neighbor 10.1.25.5, 10.1.35.5 or 10.1.45.5. You should use local preference to accomplish this task.

On R2

```
R2(config)#router bgp 100
R2(config-router)#NO bgp default local-preference 200
R2(config-router)#do cle ip bgp * out
```

On R3

```
R3(config)#router bgp 100
R3(config-router)#NO bgp default local-preference 300
R3(config-router)#do cle ip bgp * out
```

On R4

```
R4(config-router)#NO bgp default local-preference 400
R4(config-router)#do cle ip bgp * out
```

To verify the configuration:

On R1

R1#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*  i5.0.0.0         10.1.1.3                 0    100      0 200 i
*>i                 10.1.1.2                 0    100      0 200 i
*  i                10.1.1.4                 0    100      0 200 i
```

To configure the task:

On R2

```
R2(config)#route-map TST permit 10
R2(config-route-map)#set local-preference 200
R2(config-route-map)#route-map TST permit 20

R2(config)#Router bgp 100
R2(config-router)#neighbor 10.1.1.1 route-map TST OUT
R2(config-router)#do cle ip bgp * out
```

On R3

```
R3(config)#route-map TST permit 10
R3(config-route-map)#set local-preference 300
```

```
R3(config-route-map)#route-map TST permit 20

R3(config)#Router bgp 100
R3(config-router)#neighbor 10.1.1.1 route-map TST OUT
R3(config-router)#do cle ip bgp * out
```

On R4

```
R4(config)#route-map TST permit 10
R4(config-route-map)#set local-preference 200
R4(config-route-map)#route-map TST permit 20

R4(config)#Router bgp 100
R4(config-router)#neighbor 10.1.1.1 route-map TST OUT
R4(config-router)#do cle ip bgp * out
```

To test the configuration:

On R1

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*  i5.0.0.0         10.1.1.3                 0    300      0 200 i
*  i                10.1.1.2                 0    200      0 200 i
*>i                 10.1.1.4                 0    400      0 200 i

R1#Traceroute 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 10.1.1.4 0 msec 4 msec 0 msec
  2 10.1.45.5 28 msec *  28 msec
```

Task 5

Erase the startup configuration on all routers before proceeding to the next lab.

Lab 13 - The AS-Path Attribute

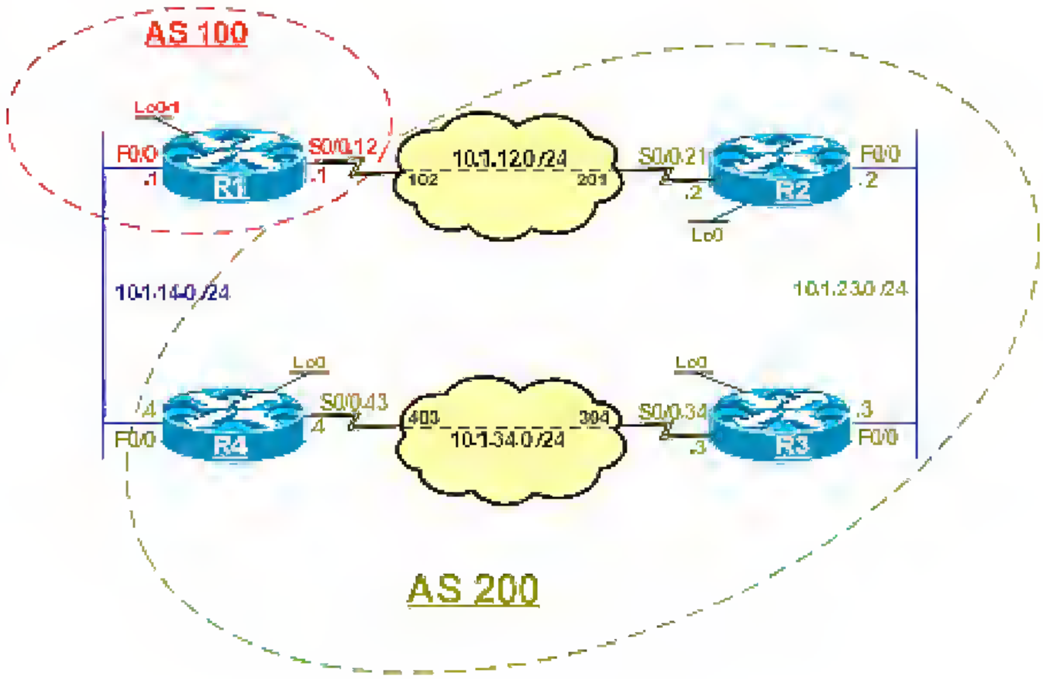

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	200
R4	Lo0 F0/0 S0/0.43	4.4.4.4 /8 10.1.14.4 /24 10.1.34.4 /24	200

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. Routers R2 and router R4 should have EBGP peer session to R1 in AS 100. BGP routers should ONLY advertise their loopback interface/s in BGP. Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1

```
R1(config-if)#router bgp 100
R1(config-router)#no au
R1(config-router)#neighbor 10.1.14.4 remote-as 200
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no au
```

```
R2(config-router)#no syn
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
R2(config-router)#neighbor 10.1.34.4 remote-as 200

R2(config)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3

```
R3(config)#router bgp 200
R3(config-router)#no au
R3(config-router)#no syn
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 200
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4

```
R4(config)#router bgp 200
R4(config-router)#no syn
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 200
R4(config-router)#neighbor 10.1.23.2 remote-as 200
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#no au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

Task 2

Configure R1 in AS 100 such that routers in AS 200 use the link through R4-R1 to reach its network 1.0.0.0 /8. Use the AS-Path attribute to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined, as followed:

On R2

R2#Sh ip bgp

```
BGP table version is 6, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric  LocPrf  Weight Path
* i1.0.0.0          10.1.14.1            0     100       0 100 i
*>                  10.1.12.1            0               0 100 i
*> 2.0.0.0          0.0.0.0              0           32768 i
*>i3.0.0.0          10.1.23.3            0     100       0 i
*>i4.0.0.0          10.1.34.4            0     100       0 i
* i11.0.0.0         10.1.14.1            0     100       0 100 i
*>                  10.1.12.1            0               0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 6, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric  LocPrf Weight Path
* i1.0.0.0          10.1.14.1            0     100      0 100 i
*>i                 10.1.12.1            0     100      0 100 i
*>i2.0.0.0          10.1.23.2            0     100      0 i
*> 3.0.0.0          0.0.0.0              0          32768 i
*>i4.0.0.0          10.1.34.4            0     100      0 i
* i11.0.0.0         10.1.14.1            0     100      0 100 i
*>i                 10.1.12.1            0     100      0 100 i
```

On R4

R4#Sh ip bgp

BGP table version is 6, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0             0 100 i
* i                 10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 11.0.0.0         10.1.14.1                0             0 100 i
* i                 10.1.12.1                0    100      0 100 i
```

To configure R1 so the routers in AS 200 take the R4-R1 link to reach network 1.0.0.0 /8:

On R1

R1(config)#**access-list 1 permit 1.0.0.0 0.255.255.255**

R1(config)#route-map TST1 permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set as-path prepend 100 100 100 100

R1(config)#route-map TST1 permit 20

R1(config-route-map)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 route-map TST1 out

To verify the configuration:

On R2

R2#Sh ip bgp

BGP table version is 7, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

```
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 100 i
*                   10.1.12.1                0             0 100 100 100 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 7, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>i                 10.1.12.1                0    100      0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 6, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0             0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 11.0.0.0         10.1.14.1                0             0 100 i
```

```
* i               10.1.12.1          0      100        0 100 i
```

Note the AS-path attribute is used to influence the degree of preference in another AS. R2, R3 and R4 will go through R4 to reach network 1.0.0.0 /8.

Task 3

Configure R1 in AS 100 such that the routers in AS 200 use the link through R2-R1 to reach network 11.0.0.0 /8. Use the AS-Path attribute to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined, as followed:

On R2

R2#Sh ip bgp

```
BGP table version is 7, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 100 i
*                   10.1.12.1                0             0 100 100 100 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 7, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
```

```
*>i1.0.0.0         10.1.14.1          0      100        0 100 i
*>i2.0.0.0         10.1.23.2          0      100        0 i
*> 3.0.0.0         0.0.0.0            0             32768 i
*>i4.0.0.0         10.1.34.4          0      100        0 i
* i11.0.0.0        10.1.14.1          0      100        0 100 i
*>i                10.1.12.1          0      100        0 100 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 6, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric  LocPrf Weight Path
*> 1.0.0.0          10.1.14.1          0                 0 100 i
*>i2.0.0.0          10.1.23.2          0      100        0 i
*>i3.0.0.0          10.1.34.3          0      100        0 i
*> 4.0.0.0          0.0.0.0            0             32768 i
*> 11.0.0.0         10.1.14.1          0                 0 100 i
* i                 10.1.12.1          0      100        0 100 i
```

To configure the attribute on R1:

On R1

R1(config)#access-list 11 permit 11.0.0.0 0.255.255.255

R1(config)#route-map **TST11** permit 10
R1(config-route-map)#match ip addr 11
R1(config-route-map)#**set as-path prepend** 100 100 100 100

R1(config)#route-map **TST11** permit 20

R1(config-route-map)#router bgp 100
R1(config-router)#neighbor 10.1.14.4 route-map **TST11** out

To verify the configuration:

On R4

R4#Show ip bgp

BGP table version is 28, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network        Next Hop      Metric  LocPrf Weight Path
*> 1.0.0.0        10.1.14.1          0              0 100 i
*>i2.0.0.0        10.1.23.2          0     100      0 i
*>i3.0.0.0        10.1.34.3          0     100      0 i
*> 4.0.0.0        0.0.0.0            0          32768 i
*  11.0.0.0       10.1.14.1          0              0 100 100 100 100 100 i
*>i               10.1.12.1          0     100      0 100 i
```

On R3

R3#Sh ip bgp

BGP table version is 13, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? – incomplete

```
   Network        Next Hop      Metric  LocPrf Weight Path
*>i1.0.0.0        10.1.14.1          0     100      0 100 i
*>i2.0.0.0        10.1.23.2          0     100      0 i
*> 3.0.0.0        0.0.0.0            0          32768 i
*>i4.0.0.0        10.1.34.4          0     100      0 i
*>i11.0.0.0       10.1.12.1          0     100      0 100 i
```

On R2

R2#Sh ip bgp

BGP table version is 13, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network        Next Hop      Metric  LocPrf Weight Path
*>i1.0.0.0        10.1.14.1          0     100      0 100 i
*                 10.1.12.1          0              0 100 100 100 100 100 i
*> 2.0.0.0        0.0.0.0            0          32768 i
```

```
*>i3.0.0.0          10.1.23.3            0       100          0 i
*>i4.0.0.0          10.1.34.4            0       100          0 i
*> 11.0.0.0         10.1.12.1            0                    0 100 i
```

Note all the routers will take the R2-R1 link to connect to network 11.0.0.0/8.

Task 4

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 14 - The Weight Attribute

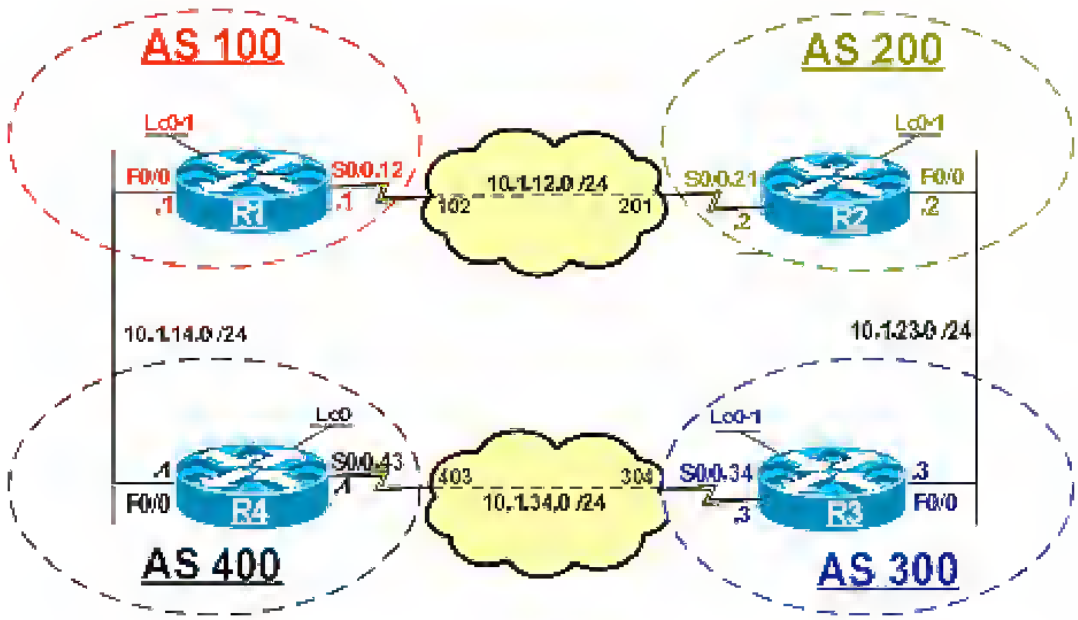

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	Lo0 Lo1 S0/0.21 F0/0	2.2.2.2 /8 22.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	Lo0 F0/0 S0/0.43	4.4.4.4 /8 10.1.14.4 /24 10.1.34.4 /24	400

Task 1

Configure router R1 in AS 100 to establish EBGP peer sessions with R2 in AS 200 and R4 in AS 400.
Router R2 should establish EBGP peer sessions with R1 in AS 100 and R3 in AS 300.
Router R3 should establish EBGP peer sessions with R2 in AS 200 and R4 in AS 400.
Router R4 should establish EBGP peer sessions with R1 in AS 100 and R3 in AS 300.
The BGP routers should ONLY advertise their loopback/s in BGP. Provide NLRI for the links using RIPv2.

On R1

```
R1(config-if)#router bgp 100
R1(config-router)#no au
R1(config-router)#neighbor 10.1.14.4 remote-as 400
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no au
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300

R2(config)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3

```
R3(config)#router bgp 300
R3(config-router)#no au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 400
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4

```
R4(config)#router bgp 400
R4(config-router)#no syn
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 300
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#no au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R4

R4#Sh ip bgp

```
BGP table version is 8, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.14.1                              0 100 200 i
*                   10.1.34.3                              0 300 200 i
*  3.0.0.0          10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 11.0.0.0         10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 22.0.0.0         10.1.14.1                              0 100 200 i
*                   10.1.34.3                              0 300 200 i
*  33.0.0.0         10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
```

Task 2

Configure R1 in AS 100 to use AS 200 to reach all the prefixes within this topology; you must use The Weight attribute to accomplish this task.

The BGP table of R1 should be examined before the weight attribute is manipulated

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

```
*   3.0.0.0       10.1.14.4                              0 400 300 i
*>                10.1.12.2                              0 200 300 i
*   4.0.0.0       10.1.12.2                              0 200 300 400 i
*>                10.1.14.4           0                  0 400 i
*> 11.0.0.0       0.0.0.0             0              32768 i
*> 22.0.0.0       10.1.12.2           0                  0 200 i
*   33.0.0.0      10.1.14.4                              0 400 300 i
*>                10.1.12.2                              0 200 300 i
```

On R1

R1(config)#router bgp 100
R1(config-router)#**neighbor 10.1.12.2 weight 40000**

The weight attribute is a Cisco-defined attribute that is local to the router. This attribute is NOT advertised to any BGP neighbor. If there is more than one route to a given destination, the weight attribute can decide which path is better. The higher the value the better the preference.

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network       Next Hop       Metric  LocPrf Weight  Path
*> 1.0.0.0        0.0.0.0           0            32768  i
*> 2.0.0.0        10.1.12.2         0            40000  200 i
*                 10.1.14.4                          0  400 300 200 i
*> 3.0.0.0        10.1.12.2                      40000  200 300 i
*                 10.1.14.4                          0  400 300 i
*> 4.0.0.0        10.1.12.2                      40000  200 300 400 i
*                 10.1.14.4         0                0  400 i
*> 11.0.0.0       0.0.0.0           0            32768  i
*> 22.0.0.0       10.1.12.2         0            40000  200 i
*                 10.1.14.4                          0  400 300 200 i
*> 33.0.0.0       10.1.12.2                      40000  200 300 i
*                 10.1.14.4                          0  400 300 i
```

Note R1 has to traverse through R2 to reach all the networks within this topology.

Task 3

The policy of AS 100 is changed to the following:

- R1 in AS 100 should use R4 in AS 400 to reach networks 33.0.0.0 /8 and network 4.0.0.0 /8.

- Network 3.0.0.0 /8 and existing and future prefixes from AS 200 should have a weight attribute of 54000 through R2.

The BGP table of R1 should be examined before implementing the weight attribute.

On R1

```
R1(config)#access-list 1 permit 33.0.0.0 0.255.255.255
R1(config)#access-list 1 permit 4.0.0.0 0.255.255.255
```

The above access-list identifies networks 33.0.0.0 /8 and 4.0.0.0 /8

```
R1(config)#access-list 2 permit 3.0.0.0 0.255.255.255
```

The above access-list identifies networks 3.0.0.0 /8

```
R1(config)#ip as-path access-list 1 permit ^200$
```

The above as-path access-list identifies existing and future prefixes in AS 200

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set weight 45000
R1(config)#route-map TST permit 20
```

The above route-map (TST) assigns a weight attribute of 45000 to the networks identified in the as-path access-list 1.

```
R1(config-route-map)#route-map TEST per 10
R1(config-route-map)#match as-path 1
R1(config-route-map)#set weight 54000
```

```
R1(config-route-map)#route-map TEST per 20
R1(config-route-map)#match ip addr 2
R1(config-route-map)#set weight 54000

R1(config-route-map)#route-map TEST per 30
```

The above route-map (TEST) assigns a weight attribute of 54000 to network 3.0.0.0 and existing and future networks advertised in AS 200.

```
R1(config-route-map)#router bgp 100
R1(config-router)#neighbor 10.1.14.4 route-map TST in
R1(config-router)#neighbor 10.1.12.2 route-map TEST in
```

The above commands assign the attributes to neighbors R2 and R4.

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0              0         32768 i
*> 2.0.0.0          10.1.12.2            0         54000 200 i
*                   10.1.14.4                          0 400 300 200 i
*> 3.0.0.0          10.1.12.2                      54000 200 300 i
*                   10.1.14.4                          0 400 300 i
*  4.0.0.0          10.1.12.2                      40000 200 300 400 i
*>                  10.1.14.4            0         45000 400 i
*> 11.0.0.0         0.0.0.0              0         32768 i
*> 22.0.0.0         10.1.12.2            0         54000 200 i
*                   10.1.14.4                          0 400 300 200 i
*  33.0.0.0         10.1.12.2                      40000 200 300 i
*>                  10.1.14.4                      45000 400 300 i
```

Note networks from AS 200 – Network 3.0.0.0 /8 all have a weight attribute of 54000

Note network 4.0.0.0 from AS 400 and network 33.0.0.0 /8 have a weight attribute of 45000.

Task 4

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 15
Multi Exit Discriminator Attribute

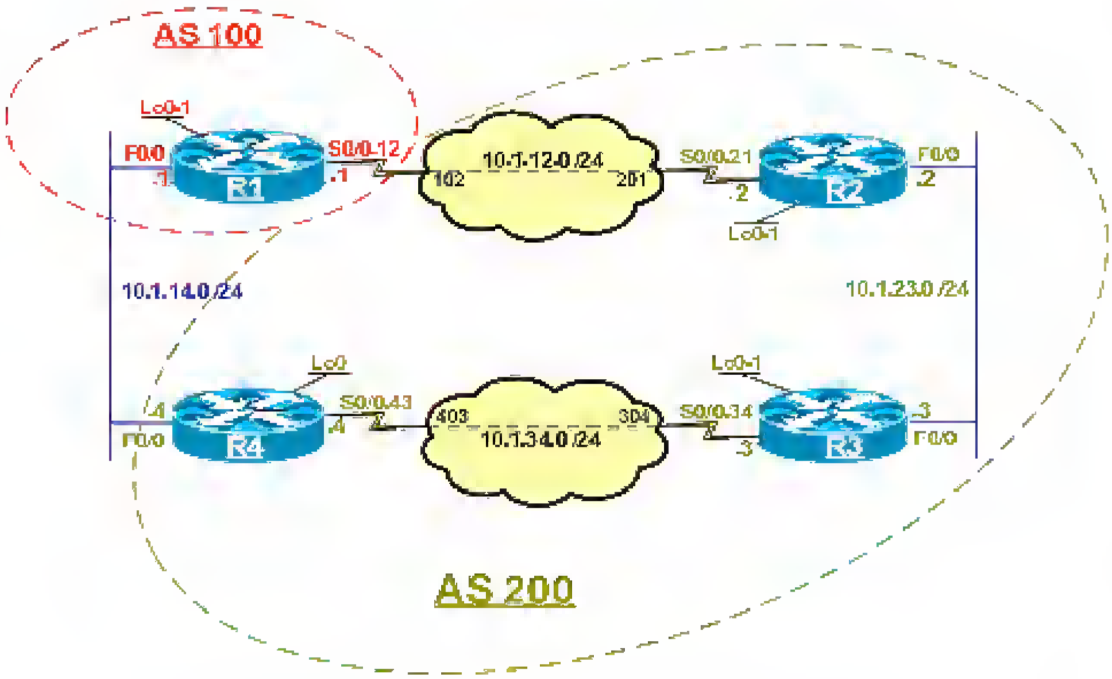

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

Ip addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	Lo0 Lo1 S0/0.21 F0/0	2.2.2.2 /8 22.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	200
R4	Lo0 F0/0 S0/0.43	4.4.4.4 /8 10.1.14.4 /24 10.1.34.4 /24	200

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. Routers R2 and router R4 should have EBGP peer session to R1 in AS 100. BGP routers should ONLY advertise their loopback interface/s in BGP. Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1

```
R1(config-if)#router bgp 100
R1(config-router)#no au
R1(config-router)#neighbor 10.1.14.4 remote-as 200
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#netw 11.0.0.0
R1(config-router)#netw 1.0.0.0

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#network 10.0.0.0
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no au
R2(config-router)#no syn
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 200
R2(config-router)#neighbor 10.1.34.4 remote-as 200

R2(config)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3

```
R3(config)#router bgp 200
R3(config-router)#no au
R3(config-router)#no syn
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.34.4 remote-as 200
R3(config-router)#neighbor 10.1.23.2 remote-as 200

R3(config)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4

```
R4(config)#router bgp 200
R4(config-router)#no syn
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 200
R4(config-router)#neighbor 10.1.23.2 remote-as 200
R4(config-router)#neighbor 10.1.14.1 remote-as 100

R4(config)#router rip
R4(config-router)#no au
R4(config-router)#ver 2
R4(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
*> 4.0.0.0          10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  33.0.0.0         10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 8, local router ID is 22.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*>i33.0.0.0         10.1.23.3                0    100      0 i
```

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
*> 4.0.0.0          10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
*  33.0.0.0         10.1.14.4                              0 200 i
*>                  10.1.12.2                              0 200 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 8, local router ID is 22.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.23.3                0    100      0 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
* i11.0.0.0         10.1.14.1                0    100      0 100 i
*>                  10.1.12.1                0             0 100 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*>i33.0.0.0         10.1.23.3                0    100      0 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 10, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.14.1                0    100      0 100 i
* i                 10.1.12.1                0    100      0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.34.4                0    100      0 i
*>i11.0.0.0         10.1.14.1                0    100      0 100 i
* i                 10.1.12.1                0    100      0 100 i
*>i22.0.0.0         10.1.23.2                0    100      0 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 8, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          10.1.12.1                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
*>i2.0.0.0          10.1.23.2                0    100      0 i
*>i3.0.0.0          10.1.34.3                0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         10.1.12.1                0    100      0 100 i
*>                  10.1.14.1                0             0 100 i
*>i22.0.0.0         10.1.23.2                0    100      0 i
*>i33.0.0.0         10.1.34.3                0    100      0 i
```

Task 2

Configure AS 200 such that router R1 in AS 100 takes R4 to reach any prefix advertised in AS 200. Manipulate MED to accomplish this task.

The following output shows the existing BGP table of R1 before manipulating the MED attribute:

On R1

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0              0         32768 i
*  2.0.0.0          10.1.14.4                          0 200 i
*>                  10.1.12.2            0             0 200 i
*  3.0.0.0          10.1.14.4                          0 200 i
*>                  10.1.12.2                          0 200 i
*> 4.0.0.0          10.1.14.4            0             0 200 i
*> 11.0.0.0         0.0.0.0              0         32768 i
*  22.0.0.0         10.1.14.4                          0 200 i
*>                  10.1.12.2            0             0 200 i
*  33.0.0.0         10.1.14.4                          0 200 i
*>                  10.1.12.2                          0 200 i
```

Note the output of the "Show ip bgp" command on R1 shows that some of the networks are reachable through R2 and some through R4

On R2

```
R2(config)#route-map TST permit 10
R2(config-route-map)#set metric 100

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST out
```

MED is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. The reason suggestion is used here is because the AS that is receiving MED attribute could use another attribute

such as Weight that will override the MED attribute. As far as MED is concerned, the lower value has the better preference.

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 12, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
*> 3.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
*  4.0.0.0          10.1.12.2              100             0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
*> 33.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2              100             0 200 i
```

Note R1 is taking R4 to reach all the networks in this topology

Task 3

Remove the configuration command from the previous task before proceeding to the next task.

On R2

R2(config)#NO route-map TST

R2(config)#router bgp 200
R2(config-router)#NO neighbor 10.1.12.1 route-map TST out

Task 4

Configure AS 200 such that AS 100 goes through R4 to reach prefix 33.0.0.0 /24 and R2 to reach Prefix 3.0.0.0 /8. Utilize MED to accomplish this task.

The following output shows the existing BGP table of R1 before manipulating the MED attribute:

On R1

R1#Sh ip bgp

```
BGP table version is 12, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.14.4                              0 200 i
*                   10.1.12.2                              0 200 i
*  4.0.0.0          10.1.12.2                              0 200 i
*>                  10.1.14.4                0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2                0             0 200 i
*> 33.0.0.0         10.1.14.4                              0 200 i
*                   10.1.12.2                              0 200 i
```

Note both networks are reachable via R4

On R2

```
R2(config)#access-list 3 permit 3.0.0.0
R2(config)#access-list 33 permit 33.0.0.0
```

Note it is very easy to remember that access-list 3 is referencing network 3.0.0.0 and access-list 33 is referencing network 33.0.0.0. If possible you should choose an access-list name or number that matches the network.

```
R2(config)#route-map TST permit 10
R2(config-route-map)#match ip addr 3
```

```
R2(config-route-map)#set metric 50

R2(config)#route-map TST permit 20
R2(config-route-map)#match ip addr 33
R2(config-route-map)#set metric 100

R2(config)#route-map TST permit 30

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST out
```

On R4

```
R4(config)#access-list 3 permit 3.0.0.0
R4(config)#access-list 33 permit 33.0.0.0

R4(config)#route-map TST permit 10
R4(config-route-map)#match ip addr 3
R4(config-route-map)#set metric 100

R4(config)#route-map TST permit 20
R4(config-route-map)#match ip addr 33
R4(config-route-map)#set metric 50

R4(config)#route-map TST permit 30

R4(config)#router bgp 200
R4(config-router)#neighbor 10.1.14.1 route-map TST out
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 16, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.14.4                              0 200 i
*>                  10.1.12.2                0             0 200 i
```

```
*  3.0.0.0      10.1.14.4       100          0 200 i
*>              10.1.12.2        50          0 200 i
*  4.0.0.0      10.1.12.2                    0 200 i
*>              10.1.14.4         0          0 200 i
*> 11.0.0.0     0.0.0.0           0      32768 i
*  22.0.0.0     10.1.14.4                    0 200 i
*>              10.1.12.2         0          0 200 i
*> 33.0.0.0     10.1.14.4        50          0 200 i
*               10.1.12.2       100          0 200 i
```

Note network 3.0.0.0 is reachable via R2 and network 33.0.0.0 is reachable via R4

Task 5

Remove BGP routing protocol and any previous BGP related command/s from the previous tasks and reconfigure them based on the chart below. These routers should ONLY advertise their loopback interface/s in BGP. The BGP peering should be established as follows:
R1 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 and R1 in AS 300 and 100 respectively.

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12 F0/0	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24 10.1.14.1 /24	100
R2	Lo0 Lo1 S0/0.21 F0/0	2.2.2.2 /8 22.2.2.0 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.0 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	Lo0 F0/0 S0/0.43	4.4.4.4 /8 10.1.14.4 /24 10.1.34.4 /24	400

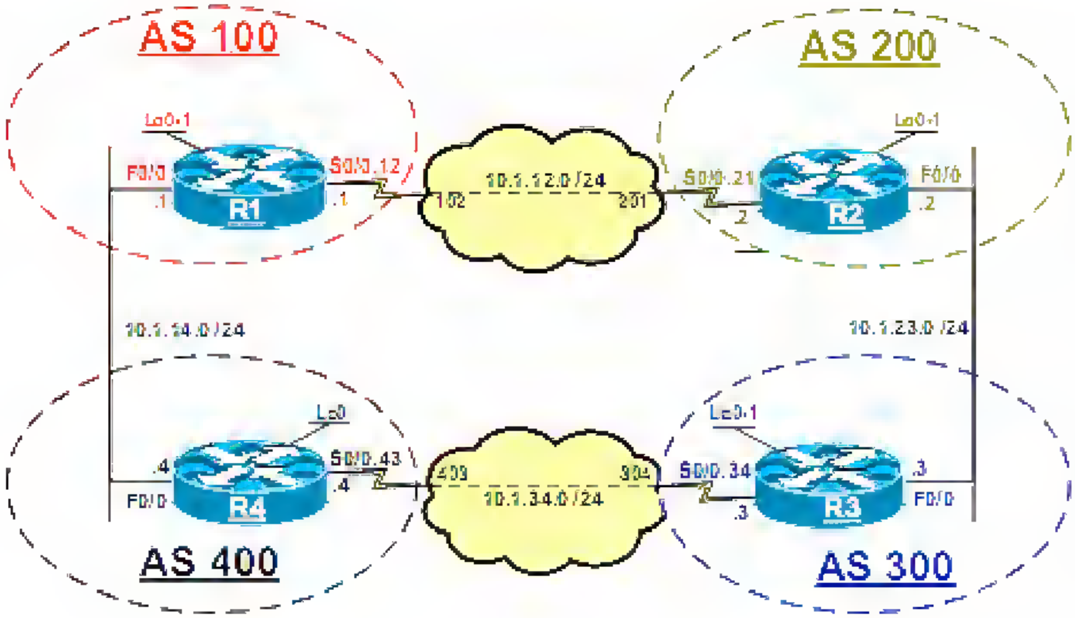

On R1

```
R1(config)#NO router bgp 100

R1(config)#router bgp 100
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 11.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.14.4 remote-as 400
```

On R2

```
R2(config)#NO router bgp 200
R2(config)#NO access-list 3
R2(config)#NO access-list 33
R2(config)#NO route-map TST

R2(config)#router bgp 200
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#no au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
```

R2(config-router)#neighbor 10.1.23.3 remote-as 300

On R3

R3(config)#NO router bgp 200

R3(config)#router bgp 300
R3(config-router)#no au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#neighbor 10.1.34.4 remote-as 400

On R4

R4(config)#NO router bgp 200
R4(config)#NO access-list 3
R4(config)#NO access-list 33
R4(config)#NO route-map TST

R4(config)#router bgp 400
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.34.3 remote-as 300

To verify the configuration:

On R1

R1#Show ip bgp

```
R1#Sh ip bgp
BGP table version is 18, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*  3.0.0.0          10.1.14.4                              0 400 300 i
*>                  10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.14.4                0             0 400 i
```

```
*> 11.0.0.0       0.0.0.0              0            32768 i
*> 22.0.0.0       10.1.12.2            0                0 200 i
*  33.0.0.0       10.1.14.4                             0 400 300 i
*>                10.1.12.2                             0 200 300 i
```

Task 6

Configure R4 in AS 400 to pass a MED of 100 to R1 in AS 100 and configure R2 in AS 200 to pass a MED of 120 to R1 in AS 100.

On R2

```
R2(config)#route-map TST permit 10
R2(config-route-map)#set metric 120

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 route-map TST out
```

On R4

```
R4(config)#route-map TST permit 10
R4(config-route-map)#set metric 100

R4(config)#router bgp 400
R4(config-router)#neighbor 10.1.14.1 route-map TST out
```

To verify the configuration:

On R1

```
R1#Show ip bgp

BGP table version is 25, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop         Metric  LocPrf Weight Path
*> 1.0.0.0        0.0.0.0               0          32768 i
```

```
*> 2.0.0.0       10.1.12.2        120                  0 200 i
*> 3.0.0.0       10.1.14.4        100                  0 400 300 i
*                10.1.12.2        120                  0 200 300 i
*> 4.0.0.0       10.1.14.4        100                  0 400 i
*> 11.0.0.0      0.0.0.0            0              32768 i
*> 22.0.0.0      10.1.12.2        120                  0 200 i
*> 33.0.0.0      10.1.14.4        100                  0 400 300 i
*                10.1.12.2        120                  0 200 300 i
```

Task 7

Ensure that R1 in AS 100 always takes R4 to reach all other prefixes advertised in this topology. You should configure R1 to accomplish this task.

On R1

R1(config)#router bgp 100
R1(config-router)#**bgp always-compare-med**
R1(config-router)#**bgp bestpath as-path ignore**

Note the tab key will not work when entering the "bgp bestpath as-path ignore" command. This command is a hidden one.

The MED as stated in RFC 1771, is an optional non-transitive attribute. The value of this attribute may be used by the BGP best path selection process to discriminate among multiple exit points to a neighboring AS. The lower the value the better the path. The MED value comparison is done only among paths from the same AS. The "bgp always-compare-med" command is used to change this behavior by enforcing MED comparison between all paths, regardless of the AS from which the paths are received.

Note the second command "bgp bestpath as-path ignore" is also needed as part of the solution in accomplishing this task, this command is needed because as-path and the origin attributes are looked at before the MED attribute. Since the Origin attribute is identical for all routes and BGP routing protocol is told to ignore the as-path attribute in best path calculation, BGP will look at the MED value next.

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
        r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop         Metric  LocPrf Weight Path
*> 1.0.0.0        0.0.0.0               0          32768 i
*> 2.0.0.0        10.1.14.4           100              0 400 300 200 i
*                 10.1.12.2           120              0 200 i
*> 3.0.0.0        10.1.14.4           100              0 400 300 i
*                 10.1.12.2           120              0 200 300 i
*> 4.0.0.0        10.1.14.4           100              0 400 i
*                 10.1.12.2           120              0 200 300 400 i
*> 11.0.0.0       0.0.0.0               0          32768 i
*> 22.0.0.0       10.1.14.4           100              0 400 300 200 i
*                 10.1.12.2           120              0 200 i
*> 33.0.0.0       10.1.14.4           100              0 400 300 i
*                 10.1.12.2           120              0 200 300 i
```

Note R1 takes R4 to reach all the networks in this topology

Task 8

Remove the BGP configuration performed in Task 6. Configure R4 in AS 400 to pass a MED value of 100 to R1 in AS 100; R2 in AS 200 should NOT pass any MED values to R1. Ensure that R1 in AS 100 takes R4 to reach any prefix advertised in this topology except the ones originated by R2. DO NOT use any global config command/s as part of the solution to this task.

To remove the configuration performed in Task 6 and verify the configuration:

On R2

```
R2(config)#router bgp 200
R2(config-router)#NO neighbor 10.1.12.1 route-map TST out

R2(config)#NO route-map TST
```

On R4

R4(config)#router bgp 400
R4(config-router)#NO neighbor 10.1.14.1 route-map TST out

R4(config)#NO route-map TST

To verify the configuration:

On R1

R1#Sh ip bgp

BGP table version is 8, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.14.4                              0 400 300 200 i
*                   10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.14.4                              0 400 300 i
*                   10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.14.4                0             0 400 i
*                   10.1.12.2                              0 200 300 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 22.0.0.0         10.1.14.4                              0 400 300 200 i
*                   10.1.12.2                0             0 200 i
*> 33.0.0.0         10.1.14.4                              0 400 300 i
*                   10.1.12.2                              0 200 300 i
```

To configure R4 to pass a MED value of 100 to R1:

On R4

R4(config)#route-map TST per 10
R4(config-route-map)#set metric 100

R4(config-route-map)#router bgp 400
R4(config-router)#neigh 10.1.14.1 route-map TST out

On R1

R1#Sh ip bgp

```
*> 2.0.0.0      10.1.12.2          0                0 200 i
*               10.1.14.4        100                0 400 300 200 i
*  3.0.0.0      10.1.12.2    4294967295             0 200 300 i
*>              10.1.14.4        100                0 400 300 i
*  4.0.0.0      10.1.12.2    4294967295             0 200 300 400 i
*>              10.1.14.4        100                0 400 i
*> 11.0.0.0     0.0.0.0            0          32768 i
*> 22.0.0.0     10.1.12.2          0                0 200 i
*               10.1.14.4        100                0 400 300 200 i
*  33.0.0.0     10.1.12.2    4294967295             0 200 300 i
*>              10.1.14.4        100                0 400 300 i
```

Note R1 takes the path through R4 to get to all the routes except the network/s advertised by R2. The reason that R1 goes through R2 to reach the networks that are originated by R2 is because R2 is giving a MED value of 0 for the networks that it originated. When R1 compares 0 to 100, it will take the route with the lowest MED.

Task 9

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 16 - Filtering Using Access-lists and Prefix-lists

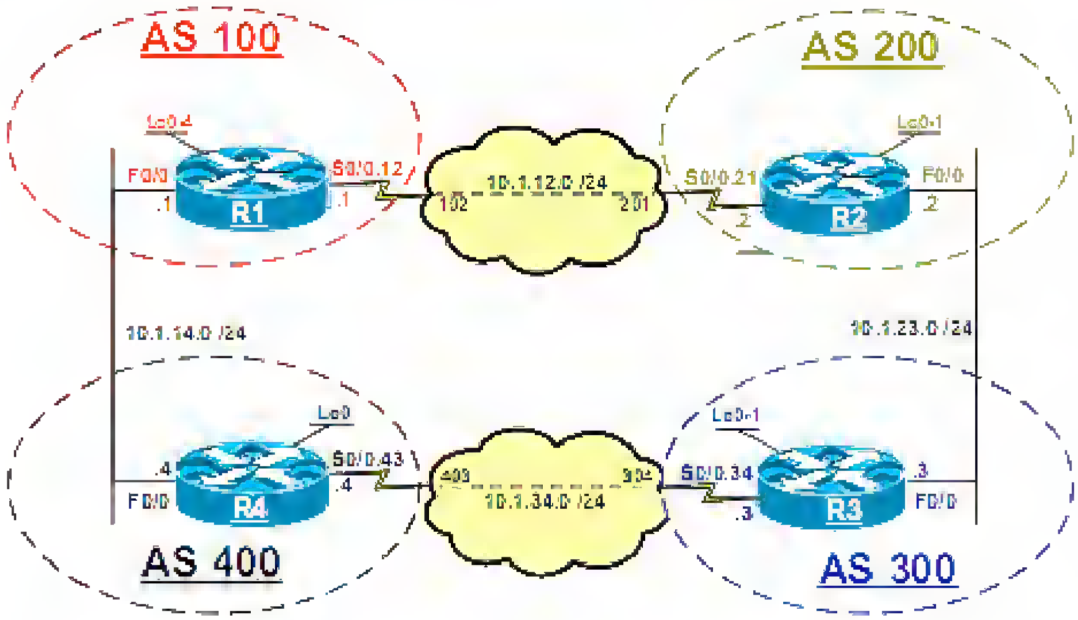

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R1 and R4's F0/0 interface should be configured in VLAN 14.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 Lo2 Lo3 Lo4 S0/0.12 F0/0	1.1.0.1 /24 1.1.1.1 /24 1.1.2.1 /24 1.1.3.1 /24 1.1.4.1 /24 10.1.12.1 /24 10.1.14.1 /24	100
R2	Lo0 Lo1 S0/0.21 F0/0	2.2.2.2 /8 22.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	Lo0 F0/0 S0/0.43	4.4.4.4 /8 10.1.14.4 /24 10.1.34.4 /24	400

Task 1

Configure the routers according to the above IP addressing chart; these routers should ONLY advertise their Loopback interface/s in BGP, the peering between the routers should be established as follows:
R1 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 and R1 in AS 300 and 100 respectively.
Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no au
R1(config-router)#netw 1.1.0.0 mask 255.255.255.0
R1(config-router)#netw 1.1.1.0 mask 255.255.255.0
R1(config-router)#netw 1.1.2.0 mask 255.255.255.0
R1(config-router)#netw 1.1.3.0 mask 255.255.255.0
R1(config-router)#netw 1.1.4.0 mask 255.255.255.0
```

```
R1(config-router)#neighbor 10.1.12.2 remote-as 200
R1(config-router)#neighbor 10.1.14.4 remote-as 400
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#no au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3

```
R3(config)#router bgp 300
R3(config-router)#no au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#neighbor 10.1.34.4 remote-as 400
```

On R4

```
R4(config)#router bgp 400
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.14.1 remote-as 100
R4(config-router)#neighbor 10.1.34.3 remote-as 300
```

To verify the configuration:

On R1

```
R1#Sh ip bgp

BGP table version is 11, local router ID is 1.1.4.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop      Metric LocPrf Weight Path
*> 1.1.0.0/24       0.0.0.0            0         32768 i
*> 1.1.1.0/24       0.0.0.0            0         32768 i
```

```
*> 1.1.2.0/24       0.0.0.0                  0             32768 i
*> 1.1.3.0/24       0.0.0.0                  0             32768 i
*> 1.1.4.0/24       0.0.0.0                  0             32768 i
*> 2.0.0.0          10.1.12.2                0                 0 200 i
*  3.0.0.0          10.1.14.4                                  0 400 300 i
*>                  10.1.12.2                                  0 200 300 i
*  4.0.0.0          10.1.12.2                                  0 200 300 400 i
*>                  10.1.14.4                0                 0 400 i
*> 22.0.0.0         10.1.12.2                0                 0 200 i
*  33.0.0.0         10.1.14.4                                  0 400 300 i
*>                  10.1.12.2                                  0 200 300 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 11, local router ID is 22.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1                0             0 100 i
*> 1.1.1.0/24       10.1.12.1                0             0 100 i
*> 1.1.2.0/24       10.1.12.1                0             0 100 i
*> 1.1.3.0/24       10.1.12.1                0             0 100 i
*> 1.1.4.0/24       10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*  4.0.0.0          10.1.12.1                              0 100 400 i
*>                  10.1.23.3                              0 300 400 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*> 33.0.0.0         10.1.23.3                0             0 300 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 11, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.0.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.1.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.2.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.3.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  1.1.4.0/24       10.1.34.4                              0 400 100 i
*>                  10.1.23.2                              0 200 100 i
*  2.0.0.0          10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*  22.0.0.0         10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 11, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.1.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.2.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.3.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 1.1.4.0/24       10.1.14.1                0             0 100 i
*                   10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.14.1                              0 100 200 i
*                   10.1.34.3                              0 300 200 i
*  3.0.0.0          10.1.14.1                              0 100 200 300 i
*>                  10.1.34.3                0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

```
*> 22.0.0.0       10.1.14.1                              0 100 200 i
*                 10.1.34.3                              0 300 200 i
*  33.0.0.0       10.1.14.1                              0 100 200 300 i
*>                10.1.34.3            0                 0 300 i
```

Task 2

Configure R2 to block network 1.1.4.0 /24 from getting into its routing and BGP tables. Use distribute-list and access-list to accomplish this task.

On R2

```
R2(config)#access-list 4 deny 1.1.4.0 0.0.0.255
R2(config)#access-list 4 permit any

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 distribute-list 4 in
R2(config-router)#neighbor 10.1.23.3 distribute-list 4 in
```

Note the tricky part was to understand the topology, if the topology in not understood, silly mistakes can occur which can cost points. The tricky part of this task is to block the prefix from both neighbors.

To verify the configuration:

On R2

R2#Sh ip bgp

```
BGP table version is 10, local router ID is 22.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop         Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1           0               0 100 i
*                   10.1.23.3                           0 300 400 100 i
*> 1.1.1.0/24       10.1.12.1           0               0 100 i
*                   10.1.23.3                           0 300 400 100 i
*> 1.1.2.0/24       10.1.12.1           0               0 100 i
*                   10.1.23.3                           0 300 400 100 i
```

```
*> 1.1.3.0/24       10.1.12.1         0             0 100 i
*                   10.1.23.3                       0 300 400 100 i
*> 2.0.0.0          0.0.0.0           0         32768 i
*  3.0.0.0          10.1.12.1                       0 100 400 300 i
*>                  10.1.23.3         0             0 300 i
*> 4.0.0.0          10.1.12.1                       0 100 400 i
*                   10.1.23.3                       0 300 400 i
*> 22.0.0.0         0.0.0.0           0         32768 i
*  33.0.0.0         10.1.12.1                       0 100 400 300 i
*>                  10.1.23.3         0             0 300 i
```

Note the network is filtered.

Task 3

Remove the configuration command from previous task, and accomplish the same task using prefix-list and distribute-list.

On R2

R2(config)#NO access-list 4

R2(config)#router bgp 200
R2(config-router)#NO neighbor 10.1.12.1 distribute-list 4 in
R2(config-router)#NO neighbor 10.1.23.3 distribute-list 4 in

R2(config)#ip prefix-list TST seq 5 deny 1.1.4.0/24
R2(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 le 32

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.12.1 prefix-list TST in
R2(config-router)#neighbor 10.1.23.3 prefix-list TST in

Note there are many ways to accomplish a given task, understanding and remembering the different ways can be the key to success.

To verify the configuration:

On R2

R2#Sh ip bgp

```
BGP table version is 10, local router ID is 22.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 1.1.1.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 1.1.2.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 1.1.3.0/24       10.1.12.1                0             0 100 i
*                   10.1.23.3                              0 300 400 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*  3.0.0.0          10.1.12.1                              0 100 400 300 i
*>                  10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.12.1                              0 100 400 i
*                   10.1.23.3                              0 300 400 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*  33.0.0.0         10.1.12.1                              0 100 400 300 i
*>                  10.1.23.3                0             0 300 i
```

Note the network is filtered.

Task 4

Configure R3 in AS 300 to block network 22.0.0.0 /8 from entering its routing and BGP table. DO NOT use distribute-list or prefix-list. A route-map and an access-list should be used to accomplish this task.

On R3

```
R3(config)#access-list 22 deny 22.0.0.0
R3(config)#access-list 22 permit any

R3(config)#route-map TST permit 10
R3(config-route-map)#match ip addr 22

R3(config)#router bgp 300
```

```
R3(config-router)#neighbor 10.1.23.2 route-map TST in
R3(config-router)#neighbor 10.1.13.1 route-map TST in
```

To verify the configuration:

On R3

R3#Sh ip bgp

```
BGP table version is 10, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.1.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.2.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.3.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.4.0/24       10.1.34.4                              0 400 100 i
*  2.0.0.0          10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*                   10.1.23.2                              0 200 100 400 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

Task 5

Remove the configuration from the previous task. Use minimum number of lines in the access-list to filter network 22.0.0.0 /8 from the BGP and IP routing table of R3. You should use an access-list and a route-map to accomplish this task.

On R3

```
R3(config)#access-list 22 permit 22.0.0.0
```

```
R3(config)#route-map TST deny 10
R3(config-route-map)#match ip addr 22

R3(config)#route-map TST permit 20

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 route-map TST in
R3(config-router)#neighbor 10.1.13.1 route-map TST in
```

To verify the configuration:

On R3

R3#Sh ip bgp

```
BGP table version is 10, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.1.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.2.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.3.0/24       10.1.34.4                              0 400 100 i
*                   10.1.23.2                              0 200 100 i
*> 1.1.4.0/24       10.1.34.4                              0 400 100 i
*  2.0.0.0          10.1.34.4                              0 400 100 200 i
*>                  10.1.23.2              0               0 200 i
*> 3.0.0.0          0.0.0.0                0           32768 i
*> 4.0.0.0          10.1.34.4              0               0 400 i
*                   10.1.23.2                              0 200 100 400 i
*> 33.0.0.0         0.0.0.0                0           32768 i
```

When we are asked to configure an access-list with minimum number of lines, we should always see if the task can be accomplished using a single statement in the access-list.

Task 6

Configure R4 to filter network 1.1.0.0 /24 – 1.1.4.0 /24 from it's BGP and IP routing table, this filtering should be configured on R4, but R1 and R3 should actually perform the filtering.

On R4

R4(config)#ip prefix-list TST seq 5 deny 1.1.0.0/24
R4(config)#ip prefix-list TST seq 10 deny 1.1.1.0/24
R4(config)#ip prefix-list TST seq 15 deny 1.1.2.0/24
R4(config)#ip prefix-list TST seq 20 deny 1.1.3.0/24
R4(config)#ip prefix-list TST seq 25 deny 1.1.4.0/24
R4(config)#ip prefix-list TST seq 30 permit 0.0.0.0/0 le 32

R4(config)#router bgp 400
R4(config-router)#neigh 10.1.14.1 prefix-list TST in
R4(config-router)#neigh 10.1.34.3 prefix-list TST in

To verify the configuration:

On R4

R4#Sh ip bgp

BGP table version is 8, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
* 2.0.0.0	10.1.34.3			0	300 200 i
*>	10.1.14.1			0	100 200 i
*> 3.0.0.0	10.1.34.3	0		0	300 i
*	10.1.14.1			0	100 200 300 i
*> 4.0.0.0	0.0.0.0	0		32768	i
*> 22.0.0.0	10.1.14.1			0	100 200 i
*> 33.0.0.0	10.1.34.3	0		0	300 i
*	10.1.14.1			0	100 200 300 i

Note even though the networks are filtered, but they are filtered on R4 and NOT on R3 or R1, to prove this, perform the following:

On R1

R1#Sh ip bgp neighbor 10.1.14.4 advertised-routes

```
BGP table version is 46, local router ID is 1.1.4.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       0.0.0.0                  0         32768 i
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*> 1.1.2.0/24       0.0.0.0                  0         32768 i
*> 1.1.3.0/24       0.0.0.0                  0         32768 i
*> 1.1.4.0/24       0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.14.4                0             0 400 i
*> 22.0.0.0         10.1.12.2                0             0 200 i
*> 33.0.0.0         10.1.12.2                              0 200 300 i
```

On R3

R3#Sh ip bgp neighbor 10.1.34.4 advertised-routes

```
BGP table version is 24, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       10.1.23.2                              0 200 100 i
*> 1.1.1.0/24       10.1.23.2                              0 200 100 i
*> 1.1.2.0/24       10.1.23.2                              0 200 100 i
*> 1.1.3.0/24       10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
*> 33.0.0.0         0.0.0.0                  0         32768 i

Total number of prefixes 8
```

Note the filtering is not performed by R3 or R1, they are still advertising the routes to R4 and R4 is performing the filtering.
In order to actually perform the filtering on R3 and R1, the "ORF" feature of BGP can be used, once the ORF capability is exchanged between the routers,

R3 and R1 will take the inbound filtering configured on R4 and they will perform it on the outbound direction.

On R3

R3(config)#router bgp 300
R3(config-router)#address-family ipv4 unicast
R3(config-router-af)#**neighbor 10.1.34.4 capability orf prefix-list receive**

On R1

R1(config)#router bgp 100
R1(config-router)#address-family ipv4 unicast
R1(config-router-af)#**neighbor 10.1.14.4 capability orf prefix-list receive**

On R4

R4(config)#router bgp 400
R4(config-router)#address-family ipv4 unicast
R4(config-router-af)#**neighbor 10.1.14.1 capability orf prefix-list send**
R4(config-router-af)#**neighbor 10.1.34.3 capability orf prefix-list send**

Note R1 and R3 are receiving the ORF, whereas, R4 is sending them. These routers can also be configured in both directions using the keyword "both", in which case they will both send and receive ORF.
When these commands are entered, the BGP peer session will be reset by the BGP process.

On R1

R1#Sh ip bgp neighbor 10.1.14.4 advertised-routes

```
BGP table version is 51, local router ID is 1.1.4.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 22.0.0.0         10.1.12.2                0             0 200 i
*> 33.0.0.0         10.1.12.2                              0 200 300 i
```

On R3

```
R3#Sh ip bgp neighbor 10.1.34.4 advertised-routes

BGP table version is 28, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

Note R1 and R3 are no longer advertising the filtered networks to R4, because filtering is performed by these routers.

Task 7

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 17 - Regular Expressions

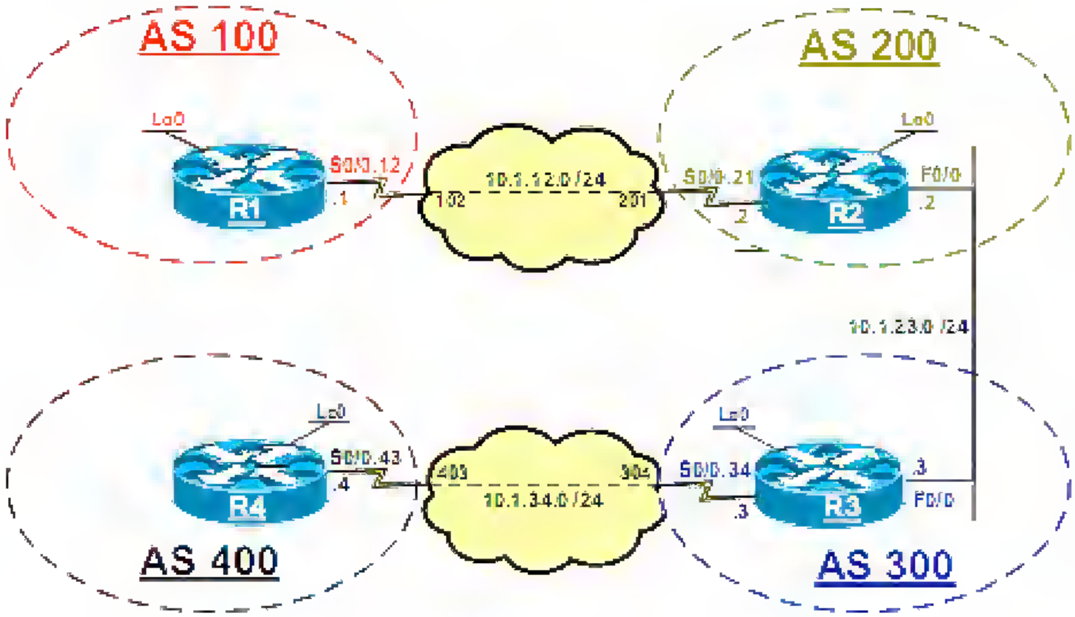

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	100
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	200
R3	Lo0 F0/0 S0/0.34	3.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	300
R4	Lo0 S0/0.43	4.4.4.4 /8 10.1.34.4 /24	400

Task 1

Configure the routers according to the above IP addressing chart; these routers should ONLY advertise their Loopback0 interface in BGP, the peering between the routers should be established as follows:
R1 should establish EBGP peer sessions with R2 in AS 200.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 in AS 300.
Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 200
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#netw 2.0.0.0
R2(config-router)#no au
R2(config-router)#neighbor 10.1.12.1 remote-as 100
R2(config-router)#neighbor 10.1.23.3 remote-as 300
```

On R3

R3(config)#router bgp 300
R3(config-router)#no au
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 200
R3(config-router)#neighbor 10.1.34.4 remote-as 400

On R4

R4(config)#router bgp 400
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 300

On All Routers:

(config)#router rip
(config-router)#no au
(config-router)#ver 2
(config-router)#Network 10.0.0.0

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 5, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
           r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1              0              0 100 i
*> 2.0.0.0          0.0.0.0                0          32768 i
*> 3.0.0.0          10.1.23.3              0              0 300 i
*> 4.0.0.0          10.1.23.3                             0 300 400 i
```

On R3

R3#Sh ip bgp

```
BGP table version is 5, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
           r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                             0 200 100 i
*> 2.0.0.0          10.1.23.2              0              0 200 i
*> 3.0.0.0          0.0.0.0                0          32768 i
*> 4.0.0.0          10.1.34.4              0              0 400 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 5, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
           r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                             0 300 200 100 i
*> 2.0.0.0          10.1.34.3                             0 300 200 i
*> 3.0.0.0          10.1.34.3              0              0 300 i
*> 4.0.0.0          0.0.0.0                0          32768 i
```

Task 2

Configure R1 such that it blocks all the prefixes that originated in AS 300.

On R1

```
R1(config)#ip as-path access-list 1 deny _300$
R1(config)#ip as-path access-list 1 permit .*

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in
```

To verify the configuration:

On R1

```
R1#Sh ip bgp

BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 3

Remove the configuration command/s from the previous task before proceeding to the next task.

On R1

```
R1(config)#NO ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#NO neighbor 10.1.12.2 filter-list 1 in
```

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 4

Configure R1 such that it blocks all the prefixes that traversed through AS 300.

On R1

```
R1(config)#ip as-path access-list 1 deny _300_
R1(config)#ip as-path access-list 1 permit .*

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in
```

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 2, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
    Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 5

Remove the configuration command from the previous task before proceeding to the next task.

On R1

R1(config)#NO ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#NO neighbor 10.1.12.2 filter-list 1 in

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 6

Configure R3 such that it doesn't advertise the prefixes that originated in it's own AS to any of it's neighbors.

On R3

R3(config)#ip as-path access-list 1 deny ^$
R3(config)#ip as-path access-list 1 permit .*

R3(config)#router bgp 300
R3(config-router)#**neighbor 10.1.23.2 filter-list 1 out**
R3(config-router)#**neighbor 10.1.34.4 filter-list 1 out**

To verify the configuration:

On R4

R4#Sh ip bgp

```
BGP table version is 18, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i
*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

On R2

R2#Sh ip bgp

```
BGP table version is 16, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

Task 7

Remove the configuration command from the previous task before proceeding to the next task.

On R3

```
R3(config)#NO ip as-path access-list 1 deny ^$

R3(config)#router bgp 300
R3(config-router)#NO neighbor 10.1.23.2 filter-list 1 out
R3(config-router)#NO neighbor 10.1.34.4 filter-list 1 out
```

To verify the configuration:

On R2

```
R2#Sh ip bgp

BGP table version is 5, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

On R3

```
R3#Sh ip bgp

BGP table version is 5, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

```
*> 4.0.0.0          10.1.34.4            0                  0 400 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 5, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                              0 300 200 100 i

*> 2.0.0.0          10.1.34.3                              0 300 200 i
*> 3.0.0.0          10.1.34.3            0                 0 300 i
*> 4.0.0.0          0.0.0.0              0             32768 i
```

Task 8

Configure R3 such that it blocks all the network from it's neighboring AS 200.

On R3

```
R3(config)#ip as-path access-list 1 deny ^200$
R3(config)#ip as-path access-list 1 permit .*

R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.23.2 filter-list 1 in
```

To verify the configuration:

On R3

R3#Sh ip bgp

```
BGP table version is 4, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0              0 400 i
```

Task 9

Remove the configuration command from the previous task before proceeding to the next task.

On R3

```
R3(config)#NO ip as-path access-list 1

R3(config)#router bgp 300
R3(config-router)#NO neighbor 10.1.23.2 filter-list 1 in
```

To verify the configuration:

On R3

```
R3#Sh ip bgp

BGP table version is 5, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0              0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0              0 400 i
```

Task 10

Configure R3 such that it blocks all the prefixes from it's directly connected neighbors.

On R3

R3(config)#ip as-path access-list 1 deny ^[0-9]+$
R3(config)#ip as-path access-list 1 permit .*

R3(config)#router bgp 300
R3(config-router)#**neighbor 10.1.23.2 filter-list 1 in**
R3(config-router)#**neighbor 10.1.34.4 filter-list 1 in**

To verify the configuration:

On R3

R3#Sh ip bgp

BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 11

Remove the configuration command from the previous task before proceeding to the next task.

On R3

R3(config)#NO ip as-path access-list 1

R3(config)#router bgp 300
R3(config-router)#NO neighbor 10.1.23.2 filter-list 1 in
R3(config-router)#NO neighbor 10.1.34.4 filter-list 1 in

To verify the configuration:

On R3

R3#Sh ip bgp

```
BGP table version is 5, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 200 100 i
*> 2.0.0.0          10.1.23.2                0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          10.1.34.4                0             0 400 i
```

Task 12

Configure R1 such that it blocks all the prefixes that originated in AS 300 and traversed through AS 200.

On R1

```
R1(config)#ip as-path access-list 1 deny _200_300$
R1(config)#ip as-path access-list 1 permit .*

R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.12.2 filter-list 1 in
```

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
```

```
*> 1.0.0.0          0.0.0.0                  0            32768 i
*> 2.0.0.0          10.1.12.2                0                0 200 i
*> 4.0.0.0          10.1.12.2                                 0 200 300 400 i
```

Task 13

Remove the configuration command from the previous task before proceeding to the next task.

On R1

```
R1(config)#NO ip as-path access-list 1

R1(config)#router bgp 100
R1(config-router)#NO neighbor 10.1.12.2 filter-list 1 in
```

To verify the configuration:

On R1

```
R1#Sh ip bgp

BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
*> 3.0.0.0          10.1.12.2                              0 200 300 i
*> 4.0.0.0          10.1.12.2                              0 200 300 400 i
```

Task 14

Your company has decided to use more complex regular expressions in the future, configure the routers such that they don't use recursive algorithm when processing regular expressions.

On R1

R1(config)#router bgp 100
R1(config-router)#**bgp regexp deterministic**

On R2

R2(config)#router bgp 200
R2(config-router)#**bgp regexp deterministic**

On R3

R3(config)#router bgp 300
R3(config-router)#**bgp regexp deterministic**

On R2

R4(config)#router bgp 400
R4(config-router)#**bgp regexp deterministic**

Task 15

Configure R1 in AS 100 to ONLY allow prefixes from its existing and future directly connected ASes, these ASes should be allowed to prepend.

Before the "As-path access-list" is written and applied, a "Show ip bgp regexp" command should be issued, if the desired output is displayed, then, the "as-path access-list" should be written and applied.
Note in the following regular expression the "(_\1)" section can be thought of the memory button in a calculator, basically the expression before it "^([0-9]+)" is what you are putting in the memory location 1, and the "*" specifies zero or more of the expression that is in the memory location 1.

On R1

R1#Sh ip bgp regexp ^([0-9]+)(_\1)*$

BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2.0.0.0	10.1.12.2	0		0	200 i

The first step is to write the "IP as-path access-list":

R1(config)#**ip as-path access-list 1 permit ^([0-9]+)(_\1)*$**

To verify the output of the "IP as-path access-list 1" command:

R1#Show ip bgp filter-list 1

BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2.0.0.0	10.1.12.2	0		0	200 i

Therefore there are two ways to test the regular expression before applying it to a neighbor/s:

1. **Show ip bgp regexp**
2. **Show ip bgp filter-list**

The next step is to apply it to neighbor/s:

R1(config)#Router bgp 100
R1(config-router)#**neighbor 10.1.12.2 filter-list 1 in**

The following uses the refresh messages so the changes can apply to the existing and new prefixes.

R1#Cle ip bgp * in

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 16

Remove the "IP as-path access-list" and the "Neighbor filter-list" commands from the previous step.

On R1

```
R1(config)#NO ip as-path access-list 1

R1(config)#Router bgp 100
R1(config-router)#NO neighbor 10.1.12.2 filter-list 1 in

R1#Clear ip bgp * in
```

Task 17

Configure R1 to prepend it's own AS number 9 times.

On R1

```
R1(config)#route-map TST permit 10
R1(config-route-map)#set as-path prepend 100 100 100 100 100 100 100 100 100

R1(config-route-map)#router bgp 200
R1(config-router)#neighbor 10.1.12.2 route-map TST out

R1#clc ip bgp * out
```

To verify the configuration:

On R1

R2#Show ip bgp

BGP table version is 21, local router ID is 2.2.2.2

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 100 100 100 100 100 100 100 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

Task 18

Configure R2 such that it allows AS-Path prepend from AS 100 ONLY if they have prepended their own AS number and NOT another AS number.

On R2

R2(config)#**ip as-path access-list 1 permit ^([0-9]+)(_\1)*$**

R2(config)#Router bgp 100
R2(config-router)#**neighbor 10.1.12.1 filter-list 1 in**

R2#Clear ip bgp * in

To verify the configuration:

On R2

R2#Sh ip bgp

```
BGP table version is 5, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
    Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.12.1                0             0 100 100 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

To test the configuration:

Change the configuration of R1 to prepend different AS numbers

On R1

R1(config)#NO route-map TST

R1(config)#route-map TST permit 10
R1(config-route-map)#**set as-path prepend 100 600 800 100**

On R2

R2#Sh ip bgp

```
BGP table version is 6, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 300 i
*> 4.0.0.0          10.1.23.3                              0 300 400 i
```

Note because AS 100 prepended other AS numbers, R2 will reject the update.

Task 19

Remove the configuration from the previous task.

On R2

R2(config)#NO ip as-path access-list 1 permit ^([0-9]+)(_\1)*$

R2(config)#router bgp 200
R2(config-router)#**NO** neighbor 10.1.12.1 filter-list 1 in

R2#Cle ip bgp *

On R1

R1(config)#**NO** route-map TST permit 10

R1(config)#router bgp 100
R1(config-router)#**NO** neighbor 10.1.12.2 route-map TST out

Task 20

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 18 - Advanced BGP configurations

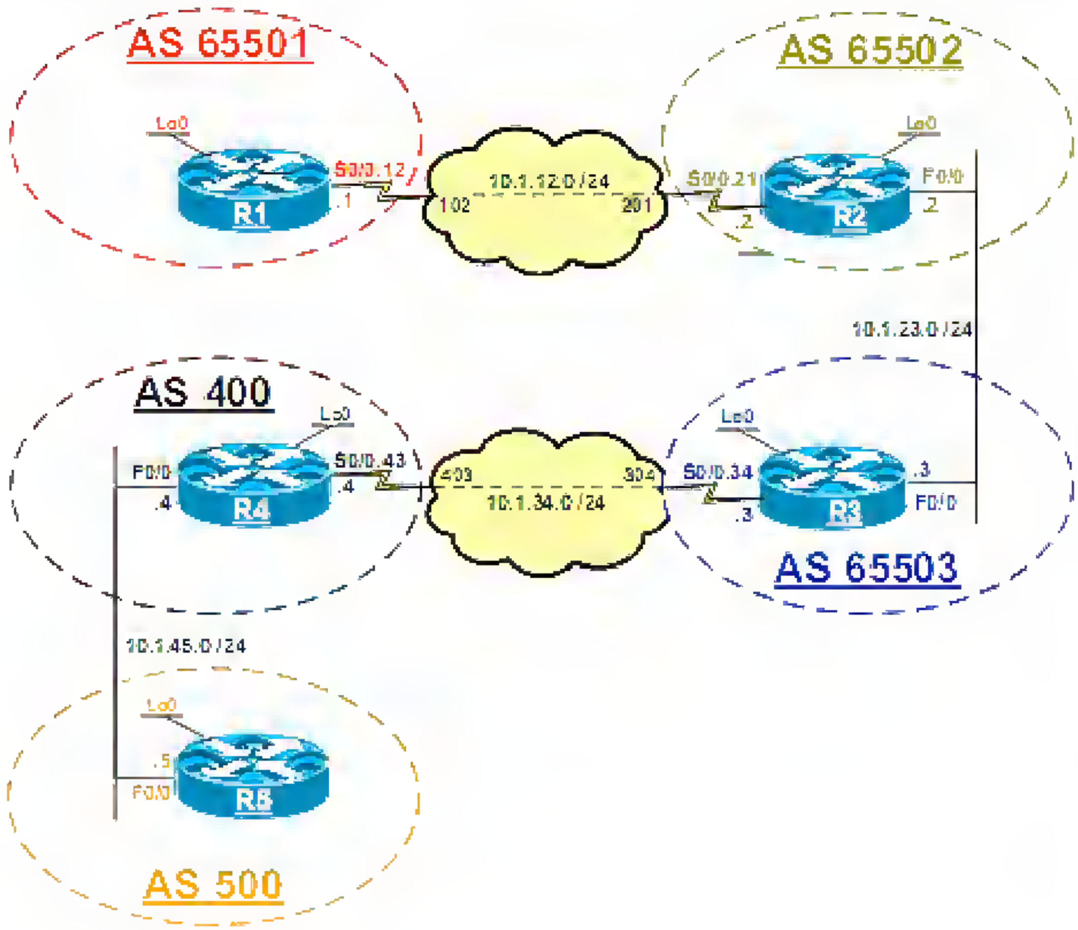

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- R4 and R5's F0/0 interface should be configured in VLAN 45.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	65501
R2	Lo0 S0/0.21 F0/0	2.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	65502
R3	Lo0 S0/0.34 F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24	65503
R4	Lo0 S0/0.43 F0/0	4.4.4.4 /8 10.1.34.4 /24 10.1.45.4 /24	400
R5	Lo0 F0/0	5.5.5.5 /8 10.1.45.5 /24	500

Task 1

Configure BGP peering on the routers as follows:

R1 in AS 65501, should establish an EBGP peer session with R2 in AS 65502.
R2 in AS 65502, should establish EBGP peer sessions with R1 and R3 in AS 65501 and 65503 respectively.
R3 in AS 65503, should establish EBGP peer sessions with R2 and R4 in AS 65502 and AS 400 respectively.
R4 in AS 400, should establish EBGP peer sessions with R3 and R5 in AS 65503 and AS 500 respectively.
R5 in AS 500 should establish an EBGP peer session with R4 in AS 400.
Provide NLRI to the links that connect the routers using RIPv2.
These routers should advertise their loopback interface/s in their assigned AS.

On R1

```
R1(config)#router bgp 65501
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 65502

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#netw 10.0.0.0
```

On R2

```
R2(config)#router bgp 65502
R2(config-router)#no au
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 65501
R2(config-router)#neighbor 10.1.23.3 remote-as 65503

R2(config-router)#router rip
R2(config-router)#ver 2
R2(config-router)#no au
R2(config-router)#netw 10.0.0.0
```

On R3

```
R3(config)#router bgp 65503
R3(config-router)#no au
R3(config-router)#netw 3.0.0.0
R3(config-router)#neighbor 10.1.23.2 remote-as 65502
R3(config-router)#neighbor 10.1.34.4 remote-as 400

R3(config-router)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
```

On R4

```
R4(config)#router bgp 400
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#neighbor 10.1.34.3 remote-as 65503
R4(config-router)#neighbor 10.1.45.5 remote-as 500

R4(config-router)#router rip
R4(config-router)#ver 2
R4(config-router)#no au
R4(config-router)#netw 10.0.0.0
```

On R5

```
R5(config-if)#router bgp 500
R5(config-router)#no au
R5(config-router)#neighbor 10.1.45.4 remote-as 400
```

```
R5(config-router)#netw 5.0.0.0

R5(config-router)#router rip
R5(config-router)#no au
R5(config-router)#ver 2
R5(config-router)#netw 10.0.0.0
```

To verify the configuration:

On R1

R1#Show ip bgp

```
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 65502 i
*> 3.0.0.0          10.1.12.2                              0 65502 65503 i
*> 4.0.0.0          10.1.12.2                              0 65502 65503 400 i
*> 5.0.0.0          10.1.12.2                              0 65502 65503 400 500 i
```

Task 2

Configure R2 such that if any of its neighbors go down, the routes from that particular neighbor are NOT removed from the BGP table. The routes should only be removed if a given neighbor is down for longer than 5 minutes.

On R2

```
R2(config-if)#router bgp 65502
R2(config-router)#NO bgp fast-external-fallover
R2(config-router)#neighbor 10.1.12.1 timers 60 300
R2(config-router)#neighbor 10.1.23.3 timers 60 300
```

On R1

R1(config-if)#router bgp 65501
R1(config-router)#neighbor 10.1.12.2 timers 60 300

On R3

R3(config-if)#router bgp 65503
R3(config-router)#neighbor 10.1.23.2 timers 60 300

This feature only supports the directly connected peers. If BGP fast external-fallover is disabled, the BGP routing process will wait until the configured hold timer expires before the peer session is reset.

Task 3

Configure R5 in AS 500 so it sets the hello and hold timer values to double its default value for its neighboring router R4.

On R5

R5(config)#router bgp 500
R5(config-router)#**neighbor 10.1.45.4 timer 120 360**

Task 4

Configure R4 such that it removes the private AS numbers when it advertises prefixes to R5.

Before configuring R4, we should display the BGP table of R5.

R5#Sh ip bgp

BGP table version is 8, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 400 65503 65502 65501 i
```

```
*> 2.0.0.0        10.1.45.4                              0 400 65503 65502 i
*> 3.0.0.0        10.1.45.4                              0 400 65503 i
*> 4.0.0.0        10.1.45.4           0                  0 400 i
*> 5.0.0.0         0.0.0.0            0              32768 i
```

On R4

```
R4(config)#router bgp 400
R4(config-router)#neighbor 10.1.45.5 remove-private-as
```

To verify the configuration:

On R5

R5#Show ip bgp

```
BGP table version is 34, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 400 i
*> 2.0.0.0          10.1.45.4                              0 400 i
*> 3.0.0.0          10.1.45.4                              0 400 i
*> 4.0.0.0          10.1.45.4                0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

Configure R1 such that it sets the minimum time between sending BGP advertisement for its EBGP neighbors to 60 seconds.

On R1

```
R1(config)#router bgp 65501
R1(config-router)#neighbor 10.1.12.2 advertisement-interval 60
```

The default values are as follows:

- External peers – 30 seconds.
- Internal peers – 5 seconds.

This command sets the minimum interval between the sending of BGP routing updates for internal and external neighbors.

Task 6

Configure the following Loopback1 interfaces on R4 and R5 and advertise them in RIPv2 only.
Ensure that R4 and R5 establish their EBGP peer session based on their Lo1 interface. Do NOT use "neighbor ebgp multihop" command to establish this peer session.
R4 - Lo1 = 44.4.4.4 /8 and R5 – Lo1 = 55.5.5.5 /8

On R4

```
R4(config)#int lo1
R4(config-if)#ip addr 44.4.4.4 255.0.0.0

R4(config-router)#router rip
R4(config-router)#netw 44.0.0.0
```

On R5

```
R5(config)#int lo1
R5(config-if)#ip addr 55.5.5.5 255.0.0.0

R5(config-router)#router rip
R5(config-router)#netw 55.0.0.0
```

On R4

```
R4(config-if)#router bgp 400
R4(config-router)#no neighbor 10.1.45.5 remote-as 500
R4(config-router)#neighbor 55.5.5.5 remote-as 500
R4(config-router)#neighbor 55.5.5.5 update-source lo1

R4(config-router)#router bgp 400
R4(config-router)#neighbor 55.5.5.5 disable-connected-check
```

On R5

```
R5(config-if)#router bgp 500
R5(config-router)#no neighbor 10.1.45.4 remote-as 400
R5(config-router)#neighbor 44.4.4.4 remote-as 400
R5(config-router)#neighbor 44.4.4.4 update-source lo1

R5(config-router)#router bgp 500
R5(config-router)#neighbor 44.4.4.4 disable-connected-check
```

The default behavior of BGP:
A BGP speaker will verify the connection of a single-hop EBGP peering session to determine if the EBGP peer is directly connected to the same network segment. If the peer is not directly connected to the same network segment, the connection verification will fail and it will prevent the peering session from being established. The "neighbor disable-connected-check" command will disable the connection verification process for EBGP peering session that are reachable by a single hop but are configured on a loopback interface or configured with a non-directly connected IP address.

Task 7

For security purposes, configure R1 and R2 such that they only accept IP packets with a TTL count in the IP header that is equal to or greater than 253. If the TTL count of the IP packet is 252 or less, these routers should ignore the packet/s.

On R1

```
R1(config)#router bgp 65501
R1(config-router)#neighbor 10.1.12.2 ttl-security hops 2
```

On R2

```
R2(config)#router bgp 65502
R2(config-router)#neighbor 10.1.12.1 ttl-security hops 2
```

This feature is enabled in 12.3(7)T and it provides a light weight security for the BGP routers against CPU utilization attacks. These types of attacks flood the network with IP packets that contain forged source and destination IP addresses in the packet headers. This configuration accepts only IP packets with a TTL count that is equal to or greater than the locally configured value.

Task 8

Remove the configuration from Tasks 6 and ensure that the two routers establish a peer session using their directly connected interface's IP address. R4 in AS 400 should establish a peer session with R5 in AS 555, and R5 in AS 500 should establish a peer session with R4 in AS 400.

On R4

R4(config)#NO int lo1

On R5

R5(config)#NO int lo1

On R4

R4(config-if)#router bgp 400
R4(config-router)#neighbor 10.1.45.5 remote-as 500
R4(config-router)#NO neighbor 55.5.5.5 remote-as 500
R4(config-router)#NO neighbor 55.5.5.5 update-source lo1

R4(config-router)#router rip
R4(config-router)#NO netw 44.0.0.0

R4(config-router)#router bgp 400
R4(config-router)#NO neighbor 55.5.5.5 **disable-connected-check**

On R5

R5(config-if)#router bgp 500
R5(config-router)#neighbor 10.1.45.4 remote-as 400
R5(config-router)#NO neighbor 44.4.4.4 remote-as 400
R5(config-router)#NO neighbor 44.4.4.4 update-source lo1

R5(config-router)#router rip
R5(config-router)#NO netw 55.0.0.0

R5(config-router)#router bgp 500
R5(config-router)#NO neighbor 44.4.4.4 **disable-connected-check**

On R4

```
R4(config-router)#NO neighbor 10.1.45.5 remote-as 500
R4(config-router)#neighbor 10.1.45.5 remote-as 555
```

On R5

```
R5(config-router)#neighbor 10.1.45.4 local-as 555
```

Typically used for AS migration and should be removed when the migration is complete. The "local-as" command prepends the AS number specified in the command to the as-path. With this command the local router appears to be in another AS.

Task 9

Configure R3 such that it replaces it's AS number with 300 and removes its private AS number when it sends updates to R4. Do NOT change the AS number of the R3 by removing the "router bgp 65503" and re-configuring "router bgp 300". Use minimum number of commands to accomplish this task.

On R3

```
R3(config)#router bgp 65503
R3(config-router)#neighbor 10.1.34.4 local-as 300 no-prepend replace-as
```

The no-prepend argument does NOT prepend the local AS number (The one configured with "router bgp" command) to the AS-Path attribute.
Replace-as argument prepends ONLY the local AS number to the AS-Path attribute that is configured after the local-as argument.

On R4

```
R4(config)#router bgp 400
R4(config-router)#NO neighbor 10.1.34.3 remote 65503
R4(config-router)#neighbor 10.1.34.3 remote-as 300
```

Task 10

Configure R3 such that it limits the number of AS-path segments that are permitted in inbound routes to 20.

On R3

R3(config)#router bgp 65503
R3(config-router)#**bgp maxas-limit 20**

This command discards routes that have a number of AS-Path segments that exceed the specified value. The range is 1 – 2000. The default value in BGP is 75. This command was introduced in IOS release 12.2.

Task 11

Configure R3 to inject a default route to neighbor 10.1.34.4 ONLY if there is a route to 2.0.0.0 /8 in R3's routing table.

On R3

R3(config)#access-list 2 permit 2.0.0.0

R3(config)#route-map **TEST** permit 10
R3(config-route-map)#match ip addr 2

R3(config)#router bgp 65503
R3(config-router)#**neighbor 10.1.34.4 default-originate route-map** TEST

Task 12

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 19 - Administrative Distance

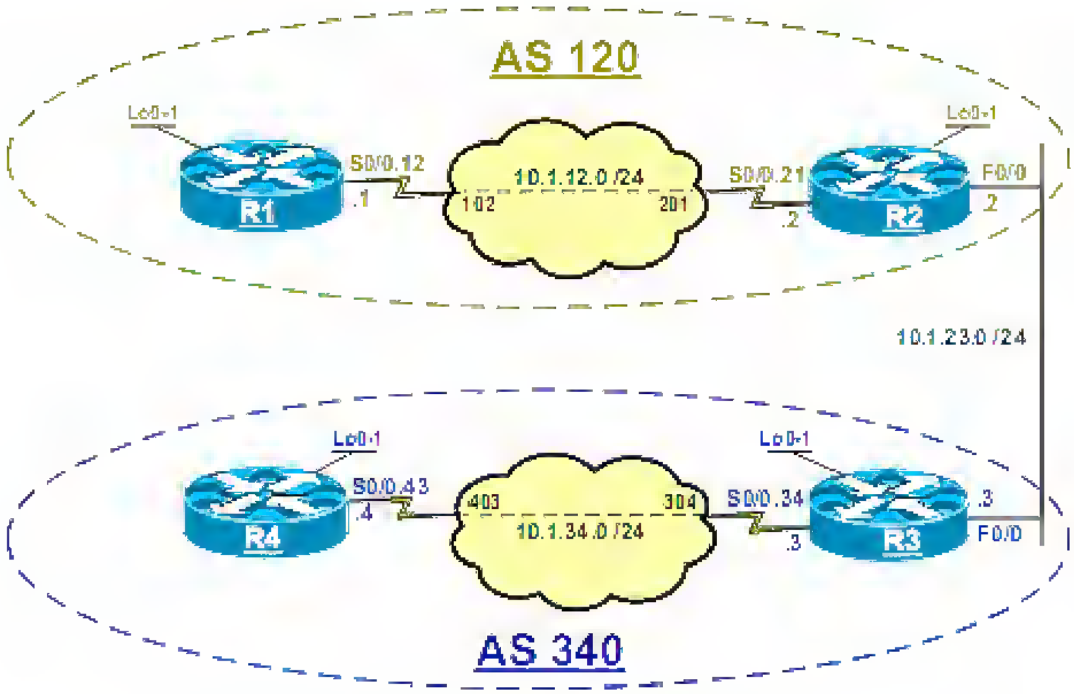

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 Lo1 S0/0.12	1.1.1.1 /8 11.1.1.1 /8 10.1.12.1 /24	120
R2	Lo0 Lo1 S0/0.21 F0/0	2.2.2.2 /8 22.2.2.2 /8 10.1.12.2 /24 10.1.23.2 /24	120
R3	Lo0 Lo1 F0/0 S0/0.34	3.3.3.3 /8 33.3.3.3 /8 10.1.23.3 /24 10.1.34.3 /24	340
R4	Lo0 Lo1 S0/0.43	4.4.4.4 /8 44.4.4.4 /8 10.1.34.4 /24	340

Task 1

Configure the routers according to the above IP addressing chart; these routers should ONLY advertise their Loopback interfaces in BGP, the peering between the routers should be established as follows:
R1 in AS 120 should establish an IBGP peer session with R2 in AS 120.
R2 should establish an IBGP peer session with R1 in AS 120 and an EBGP peer session with R3 in AS 340.
R3 should establish an EBGP peer session with R2 in AS 120 and an IBGP peer session with R4 in AS 340.
R4 should establish an IBGP peer sessions with R3 in AS 340.
Provide NLRI for the links using RIPv2, disable automatic summarization.

On R1

```
R1(config)#router bgp 120
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 11.0.0.0

R1(config-router)#neighbor 10.1.12.2 remote-as 120
```

On R2

```
R2(config)#router bgp 120
R2(config-router)#netw 2.0.0.0
R2(config-router)#netw 22.0.0.0
R2(config-router)#no au

R2(config-router)#neighbor 10.1.12.1 remote-as 120
R2(config-router)#neighbor 10.1.23.3 remote-as 340
```

On R3

```
R3(config)#router bgp 340
R3(config-router)#no au
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 33.0.0.0

R3(config-router)#neighbor 10.1.23.2 remote-as 120
R3(config-router)#neighbor 10.1.34.4 remote-as 340
```

On R4

```
R4(config)#router bgp 340
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0
R4(config-router)#netw 44.0.0.0

R4(config-router)#neighbor 10.1.34.3 remote-as 340
```

On All Routers:

```
(config)#router rip
(config-router)#no au
(config-router)#ver 2
(config-router)#Network 10.0.0.0
```

To verify the configuration:

On R1

R1#Sh ip bgp

```
BGP table version is 9, local router ID is 11.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          10.1.12.2                0    100      0 i
*>i3.0.0.0          10.1.23.3                0    100      0 340 i
*>i4.0.0.0          10.1.23.3                0    100      0 340 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*>i22.0.0.0         10.1.12.2                0    100      0 i
*>i33.0.0.0         10.1.23.3                0    100      0 340 i
*>i44.0.0.0         10.1.23.3                0    100      0 340 i
```

<u>On R2</u>

<u>R2#Sh ip bgp</u>

```
BGP table version is 9, local router ID is 22.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          10.1.12.1                0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          10.1.23.3                0             0 340 i
*> 4.0.0.0          10.1.23.3                              0 340 i
*>i11.0.0.0         10.1.12.1                0    100      0 i
*> 22.0.0.0         0.0.0.0                  0         32768 i
*> 33.0.0.0         10.1.23.3                0             0 340 i
*> 44.0.0.0         10.1.23.3                              0 340 i
```

<u>On R3</u>

<u>R3#Sh ip bgp</u>

```
BGP table version is 9, local router ID is 33.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.23.2                              0 120 i
*> 2.0.0.0          10.1.23.2                0             0 120 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

```
*>i4.0.0.0        10.1.34.4            0      100          0 i
*> 11.0.0.0       10.1.23.2                                0 120 i
*> 22.0.0.0       10.1.23.2            0                   0 120 i
*> 33.0.0.0        0.0.0.0             0              32768 i
*>i44.0.0.0       10.1.34.4            0      100          0 i
```

On R4

R4#Sh ip bgp

```
BGP table version is 9, local router ID is 44.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0        10.1.23.2            0      100          0 120 i
*>i2.0.0.0        10.1.23.2            0      100          0 120 i
*>i3.0.0.0        10.1.34.3            0      100          0 i
*> 4.0.0.0         0.0.0.0             0              32768 i
*>i11.0.0.0       10.1.23.2            0      100          0 120 i
*>i22.0.0.0       10.1.23.2            0      100          0 120 i
*>i33.0.0.0       10.1.34.3            0      100          0 i
*> 44.0.0.0        0.0.0.0             0              32768 i
```

Task 2

Configure R2 such that it changes the administrative distance of all prefixes received from R3 to 150.

On R2

```
R2(config)#router bgp 120
R2(config-router)#distance 150 10.1.23.3 0.0.0.0
```

To verify the configuration:

On R2

R2#Show ip route bgp

```
B    1.0.0.0/8 [200/0] via 10.1.12.1, 00:00:56
B    33.0.0.0/8 [150/0] via 10.1.23.3, 00:00:56
B    3.0.0.0/8 [150/0] via 10.1.23.3, 00:00:56
B    4.0.0.0/8 [150/0] via 10.1.23.3, 00:00:56
B    11.0.0.0/8 [200/0] via 10.1.12.1, 00:00:56
B    44.0.0.0/8 [150/0] via 10.1.23.3, 00:00:56
```

Task 3

Remove the configuration commands from the previous step before proceeding to the next task.

On R2

```
R2(config)#router bgp 120
R2(config-router)#NO distance 150 10.1.23.3 0.0.0.0
```

To verify the configuration

On R2

R2#Sh ip route bgp

```
B    1.0.0.0/8 [200/0] via 10.1.12.1, 00:01:02
B    33.0.0.0/8 [20/0] via 10.1.23.3, 00:01:02
B    3.0.0.0/8 [20/0] via 10.1.23.3, 00:01:02
B    4.0.0.0/8 [20/0] via 10.1.23.3, 00:01:02
B    11.0.0.0/8 [200/0] via 10.1.12.1, 00:01:02
B    44.0.0.0/8 [20/0] via 10.1.23.3, 00:01:02
```

Task 4

Configure R2 such that it ONLY changes the administrative distance of prefix 33.0.0.0 /8 to 150.

On R2

R2(config)#access-list 33 permit 33.0.0.0

R2(config)#router bgp 120
R2(config-router)#distance 150 10.1.23.3 0.0.0.0 33

Note the distance command changes the administrative distance to 150 from neighbor 10.1.23.3 followed by the inverse-mask of the IP address of the neighbor for what ever prefix that's specified in access-list 33.

To verify the configuration:

On R2

R2#Sh ip route bgp

```
B    1.0.0.0/8 [200/0] via 10.1.12.1, 00:01:36
B    33.0.0.0/8 [150/0] via 10.1.23.3, 00:01:36
B    3.0.0.0/8 [20/0] via 10.1.23.3, 00:01:36
B    4.0.0.0/8 [20/0] via 10.1.23.3, 00:01:36
B    11.0.0.0/8 [200/0] via 10.1.12.1, 00:01:36
B    44.0.0.0/8 [20/0] via 10.1.23.3, 00:01:36
```

Task 5

Remove the configuration commands from the previous step before proceeding to the next task

On R2

R2(config)#NO access-list 1 permit 33.0.0.0 0.255.255.255

R2(config)#router bgp 120
R2(config-router)#NO distance 150 10.1.23.3 0.0.0.0 1

To verify the configuration:

On R2

R2#Sh ip route bgp

```
B    1.0.0.0/8 [200/0] via 10.1.12.1, 00:00:42
B    33.0.0.0/8 [20/0] via 10.1.23.3, 00:00:42
B    3.0.0.0/8 [20/0] via 10.1.23.3, 00:00:42
B    4.0.0.0/8 [20/0] via 10.1.23.3, 00:00:42
B    11.0.0.0/8 [200/0] via 10.1.12.1, 00:00:42
B    44.0.0.0/8 [20/0] via 10.1.23.3, 00:00:42
```

Task 6

Using minimum number of commands change the administrative distance of all IBGP prefixes to 90, EBGP prefixes to 60 and locally generated prefixes to 20. This should be performed on all routers.

On R1 and R2

```
R1(config)#router bgp 120
R1(config-router)#distance bgp 60 90 20
```

On R3 and R4

```
(config)#router bgp 340
(config-router)#distance bgp 60 90 20
```

On R2

R2#Show ip route bgp

```
B    1.0.0.0/8 [90/0] via  10.1.12.1, 00:03:21
B    33.0.0.0/8 [60/0] via 10.1.23.3, 00:03:21
B      3.0.0.0/8 [60/0] via 10.1.23.3, 00:03:21
B      4.0.0.0/8 [60/0] via 10.1.23.3, 00:03:21
B    11.0.0.0/8 [90/0] via 10.1.12.1, 00:03:21
B    44.0.0.0/8 [60/0] via 10.1.23.3, 00:03:21
```

Note These are IBGP routes

Note these are EBGP routes

To see the locally generated routes, we must configure an aggregate route.

Task 7

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 20 - BGP Confederation

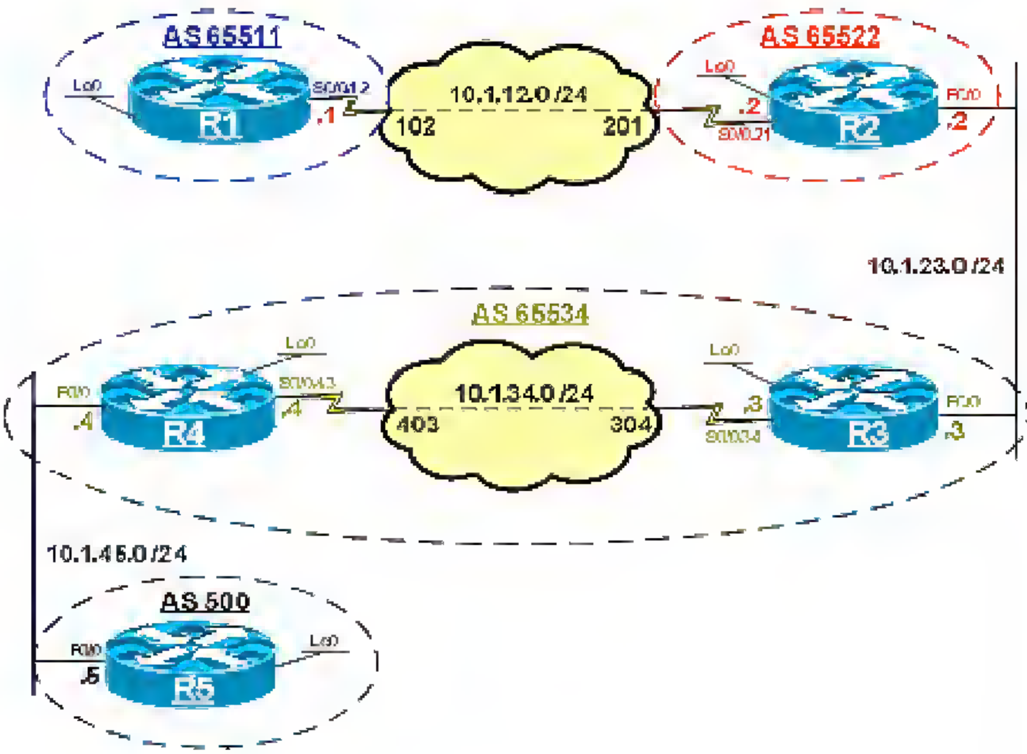

Lab Setup:

- Configure the routers that are connected to the frame-relay clouds in a point-to-point manner.
- R2 and R3's F0/0 interface should be configured in VLAN 23.
- R4 and R5's F0/0 interface should be configured in VLAN 45.
- Use the following IP addressing chart for IP address assignment.

IP addressing:

Router	Interface	IP Address	AS number
R1	Lo0 S0/0.12	1.1.1.1 /8 10.1.12.1 /24	65511
R2	**Lo0** **S0/0.21** **F0/0**	**2.2.2.2 /8** **10.1.12.2 /24** **10.1.23.2 /24**	**65522**
R3	Lo0 S0/0.34 F0/0	3.3.3.3 /8 10.1.34.3 /24 10.1.23.3 /24	65534
R4	**Lo0** **S0/0.43** **F0/0**	**4.4.4.4 /8** **10.1.34.4 /24** **10.1.45.4 /24**	**65534**
R5	Lo0 F0/0	5.5.5.5 /8 10.1.45.5 /24	500

Task 1

Configure BGP peering on the routers as follows:

- R1 in AS 65511, should establish an EBGP peer session with R2 in AS 65522.
- R2 in AS 65522, should establish EBGP peer sessions with R1 and R3 in AS 65511 and 65534 respectively.
- R3 in AS 65534, should establish an EBGP peer sessions with R2 in AS 65522 and an IBGP peer session with R4 in AS 65534.
- R4 in AS 65534, should establish an IBGP peer sessions with R3 in AS 65534 and an EBGP peer session with R5 in AS 500.
- R5 in AS 500 should establish an EBGP peer session with R4 in AS 100.
- Provide NLRI to the links that connect the routers using RIPv2.
- These routers should advertise their Loopback interface in BGP.

On All Routers

(config-router)#router rip
(config-router)#no au
(config-router)#ver 2
(config-router)#network 10.0.0.0

On R1

R1(config)#router bgp 65511
R1(config-router)#no au
R1(config-router)#netw 1.0.0.0
R1(config-router)#neighbor 10.1.12.2 remote-as 65522

R1(config-router)#**bgp confederation identifier** 100

The "bgp confederation identifier" command is used to configure a single AS number to identify a group of smaller ASes as a single confederation.
A confederation can be used to reduce the IBGP mesh by dividing a large single AS into multiple Sub-ASes and then grouping them into a single confederation.

R1(config-router)#**bgp confederation peers** 65522

The above command is used to configure multiple ASes as a single confederation. The ASes specified in this command are visible internally to the confederation.

On R2

R2(config)#router bgp 65522
R2(config-router)#no au
R2(config-router)#netw 2.0.0.0
R2(config-router)#neighbor 10.1.12.1 remote-as 65511
R2(config-router)#neighbor 10.1.23.3 remote-as 65534

R2(config-router)#**bgp confederation identifier** 100
R2(config-router)#**bgp confederation peers** 65511 65534

R2(config-router)#neighbor 131.1.23.3 next-hop-self

On R3

R3(config)#router bgp 65534

R3(config-router)#no au
R3(config-router)#no syn
R3(config-router)#neighbor 10.1.34.4 remote-as 65534
R3(config-router)#neighbor 10.1.23.2 remote-as 65522
R3(config-router)#netw 3.0.0.0

R3(config-router)#**bgp confederation identifier** 100
R3(config-router)#**bgp confederation peers** 65522

On R4

R4(config)#router bgp 65534
R4(config-router)#no syn
R4(config-router)#no au
R4(config-router)#netw 4.0.0.0

R4(config-router)#**bgp confederation identifier** 100

R4(config-router)#neighbor 10.1.45.5 remote-as 500
R4(config-router)#neighbor 10.1.34.3 remote-as 65534

On R5

R5(config)#router bgp 500
R5(config-router)#no au
R5(config-router)#netw 5.0.0.0

R5(config-router)#neighbor 10.1.45.4 remote-as 100

To verify the configuration:

On R5

R5#Show ip bgp

```
BGP table version is 6, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric  LocPrf  Weight Path
*> 1.0.0.0          10.1.45.4                            0 100 i
*> 2.0.0.0          10.1.45.4                            0 100 i
*> 3.0.0.0          10.1.45.4                            0 100 i
*> 4.0.0.0          10.1.45.4            0               0 100 i
*> 5.0.0.0          0.0.0.0              0           32768 i
```

Note to R5 all the prefixes are from AS 100.

On R1

R1#Show ip bgp

```
BGP table version is 6, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
          r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf   Weight Path
*> 1.0.0.0          0.0.0.0                  0           32768 i
*> 2.0.0.0          10.1.12.2                0    100        0 (65522) i
*> 3.0.0.0          10.1.23.3                0    100        0 (65522 65534) i
*> 4.0.0.0          10.1.34.4                0    100        0 (65522 65534) i
*> 5.0.0.0          10.1.45.5                0    100        0 (65522 65534) 500 i
```

Note the AS-Path/s in the parenthesis are the private AS numbers within the confederation, AS 500 is outside of the parenthesis because its NOT part of the confederation.

Task 2

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 21 – BGP Hiding Local AS number

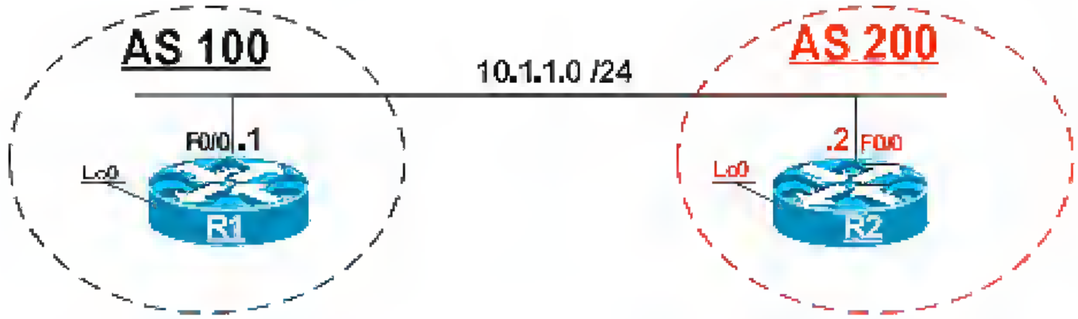

Lab Setup:

- Configure the F0/0 interface of both routers in VLAN 100
- Use the IP addressing chart for IP address assignment

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.1.1 /24 Loopback0 = 1.1.1.1 /8
R2	F0/0 = 10.1.1.2 /24 Loopback0 = 2.2.2.2 /8

Task 1

Configure R1 in AS 100 to establish an EBGP session with R2 in AS 200.

On R1

```
R1(config)#router bgp 100
R1(config-router)#no auto-summary
R1(config-router)#network 1.0.0.0

R1(config-router)#neighbor 10.1.1.2 remote-as 200
```

On R2

```
R2(config)#router bgp 200
R2(config-router)#no auto-summary
R2(config-router)#network 2.0.0.0

R2(config-router)#neighbor 10.1.1.1 remote-as 100
```

To verify the configuration:

On R1

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.1.2                 0             0 200 i
```

On R2

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note from R2's perspective prefix 1.0.0.0/8 was originated & advertised by AS 100, and from R1's perspective, prefix 2.0.0.0/8 was originated and advertised by AS 200.

Task 2

Configure R1 in AS 111 to establish an EBGP session with R2 in AS 200 such that the output of the "Show ip bgp" command on these two routers will be identical to the follows:

On R1

R1#Sh ip bgp

```
BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight  Path
*> 1.0.0.0          0.0.0.0                  0         32768  i
*> 2.0.0.0          10.1.1.2                 0             0  100 200 i
```

On R2

R2#Show ip bgp

```
BGP table version is 5, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale

   Network          Next Hop            Metric LocPrf Weight  Path
*> 1.0.0.0          10.1.1.1                 0             0  100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768  i
```

On R1

```
R1(config)#router bgp 111
R1(config-router)#no auto-summary
R1(config-router)#network 1.0.0.0

R1(config-router)#neighbor 10.1.1.2 remote-as 200
R1(config-router)#neighbor 10.1.1.2 local-as 100
```

By changing the AS number of R1 to 111, and a "Neighbor 10.1.1.2 local-as 100" R2 will see R1's real AS of 111 originating the route and then AS 100 was the AS that advertised it. Note on R1, it shows that prefix 2.0.0.0/8 was originated by AS 200 but the advertising AS to R1 was AS 100. They both see

this invisible AS 100.

Task 3

Configure R1 such that when R2 advertises network 2.0.0.0/8, the output of the "Show ip bgp" command on R1 is identical to the following:

On R1

R1#Sh ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight  Path
*> 1.0.0.0          0.0.0.0                  0         32768  i
*> 2.0.0.0          10.1.1.2                 0             0  200 i
```

On R1

```
R1(config)#router bgp 111
R1(config-router)#neighbor 10.1.1.2 local-as 100 no-prepend
```

Note the "no-prepend" option tells the router NOT to prepend AS 100 to the advertised prefixes. This will ONLY affect R1's BGP table.

On R1

R1#Show ip bgp

```
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight  Path
*> 1.0.0.0          0.0.0.0                  0         32768  i
*> 2.0.0.0          10.1.1.2                 0             0  200 i
```

On R2

R2#Show ip bgp

```
BGP table version is 7, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Task 4

Configure R1 such that the output of the "Show ip bgp" command on R2 is identical to the following:

On R2

R2#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1

```
R1(config)#router bgp 111
R1(config-router)#neighbor 10.1.1.2 local-as 100 no-prepend replace-as
```

Note the "replace-as" option instructs the local router NOT to prepend the real AS number.

On R2

R2#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

Configure R1 such that R2 can establish an EBGP peer session with R1 using AS 111 or 100.

On R1

R1(config)#router bgp 111
R1(config-router)#**neighbor 10.1.1.2 local-as 100 no-prepend replace-as dual-as**

On R1

R1#Show ip bgp

BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	0.0.0.0	0		32768	i
*> 2.0.0.0	10.1.1.2	0		0	200 i

On R2

R2#Show ip bgp

BGP table version is 13, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	10.1.1.1	0		0	100 i
*> 2.0.0.0	0.0.0.0	0		32768	i

Note the current configuration of R2 is as follows:

On R2

R2#Show run | b router bgp 200

```
router bgp 200
 no synchronization
 bgp log-neighbor-changes
 network 2.0.0.0
 neighbor 10.1.1.1 remote-as 100
 no auto-summary
```

Note its establishing a peer session with R1 using AS 100, the following verifies that R2 can also establish a peer session with R1 using AS 111:

On R2

```
R2(config)#router bgp 200
R2(config-router)#no auto-summary
R2(config-router)#netw 2.0.0.0

R2(config-router)#neighbor 10.1.1.1 remote-as 111
```

To verify the configuration:

On R2

```
R2#Show ip bgp

BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
            r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight  Path
*> 1.0.0.0          10.1.1.1                 0             0  111 i
*> 2.0.0.0          0.0.0.0                  0         32768  i
```

Note the ONLY difference here is the AS that originated and advertised 1.0.0.0/8 prefix.

Task 6

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 22 – BGP Allow-as

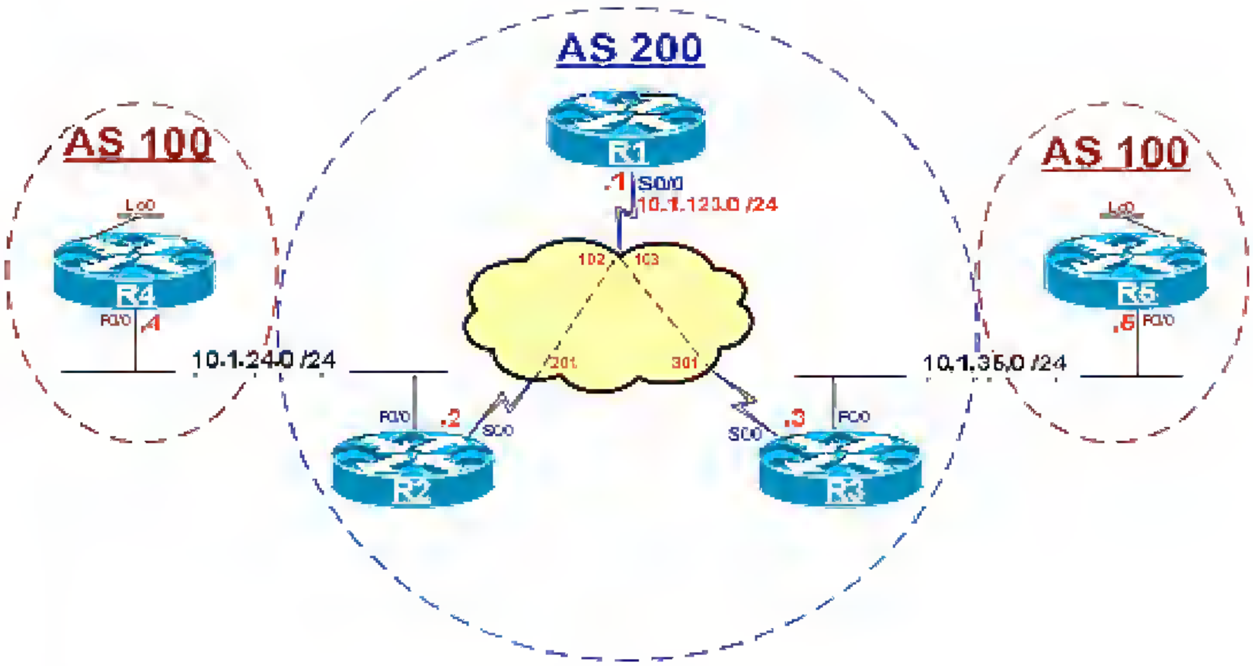

Lab Setup:

- Configure Frame-relay on R1, R2 and R3; this configuration should be done directly under the main interface. R2 and R3 should have a static Frame-relay mapping to R1 and R1 should have static frame-relay mapping to R2 and R3.
- RIPv2 should be used to provide NLRI
- Configure the F0/0 interface of R2 and R4 in VLAN 24
- Configure the F0/0 interface R3 and R5 in VLAN 35
- Use the following chart for IP addressing and AS assignment.

IP addressing & AS assignment:

Router	Interface / IP address	AS number
R1	S0/0 = 10.1.123.1 /24	200
R2	S0/0 = 10.1.123.2 /24 F0/0 = 10.1.24.2 /24	200
R3	S0/0 = 10.1.123.3 /24 F0/0 = 10.1.35.3 /24	200
R4	F0/0 = 10.1.24.4 /24 Lo0 = 4.4.4.4 /8	100
R5	F0/0 = 10.1.35.5 /24 Lo0 = 5.5.5.5 /8	100

Task 1

Configure the routers according to the above diagram; if this configuration is performed successfully all routers should have network 4.0.0.0 /8 and 5.0.0.0 /8 in their BGP and routing table.

On R1

R1(config)#router bgp 200
R1(config-router)#neighbor 10.1.123.2 remote-as 200
R1(config-router)#neighbor 10.1.123.3 remote-as 200

On R2

R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.123.1 remote-as 200
R2(config-router)#neighbor 10.1.123.3 remote-as 200
R2(config-router)#neighbor 10.1.24.4 remote-as 100

On R3

R3(config)#router bgp 200
R3(config-router)#neighbor 10.1.123.1 remote-as 200
R3(config-router)#neighbor 10.1.123.2 remote-as 200
R3(config-router)#neighbor 10.1.35.5 remote-as 100

On R4

R4(config)#router bgp 100

R4(config-router)#neighbor 10.1.24.2 remote-as 200
R4(config-router)#network 4.0.0.0

On R5

R5(config)#router bgp 100
R5(config-router)#neighbor 10.1.35.3 remote-as 200
R5(config-router)#network 5.0.0.0

To verify the configuration:

On R1

R1#Show ip bgp | b Network

```
 Network          Next Hop            Metric LocPrf Weight Path
*>i4.0.0.0        10.1.24.4                0    100      0 100 i
*>i5.0.0.0        10.1.35.5                0    100      0 100 i
```

On R2

R2#Show ip bgp | b Network

```
 Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0        10.1.24.4                0             0 100 i
*>i5.0.0.0        10.1.35.5                0    100      0 100 i
```

On R3

R3#Show ip bgp | b Network

```
 Network          Next Hop            Metric LocPrf Weight Path
*>i4.0.0.0        10.1.24.4                0    100      0 100 i
*> 5.0.0.0        10.1.35.5                0             0 100 i
```

On R4

R4#Show ip bgp | b Network

```
 Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0        0.0.0.0                  0         32768 i
```

Note R4 does NOT have prefix 5.0.0.0 /8 in its routing table, the question is did

R2 advertise that prefix to R4? The output of the following show command can be used to verify:

On R2

R2#Show ip bgp neighbor 10.1.24.4 advertised-routes | b Network

```
   Network          Next Hop            Metric  LocPrf  Weight Path
*>i5.0.0.0          10.1.35.5                0     100       0 100 i
```

Total number of prefixes 1

Note R2 has advertised prefix 5.0.0.0 /8 but this prefix came from AS 100 and since R4 is in the same AS (AS 100), it will discard that prefix, this is done by BGP as loop avoidance.
To fix this problem, R4 should be configured to allow prefix/es that have its own AS number in its AS, this is accomplished using the following configuration:

On R4

R4(config)#router bgp 100
R4(config-router)#**neighbor 10.1.24.2** allowas-in

To verify the configuration:

On R4

R4#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.24.2                              0 200 100 i
```

R4#Show ip route | Inc B

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
B   5.0.0.0/8 [20/0] via 10.1.24.2, 00:16:20

Note prefix 5.0.0.0 /8 is in R4's BGP and routing table with AS number that matches its own.
Since R5 has the same problem, it should also be configured the same way.

To verify R5's BGP table before configuration:

On R5

R5#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Note once again prefix 4.0.0.0 /8 is missing.

To fix the problem:

On R5

R5(config)#router bgp 100
R5(config-router)#**neighbor 10.1.35.3 allowas-in**

To verify the configuration:

On R5

R5#Show ip bgp | b Network

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          10.1.35.3                              0 200 100 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

Policy Based Routing

Lab 1 – PBR based on Source IP address

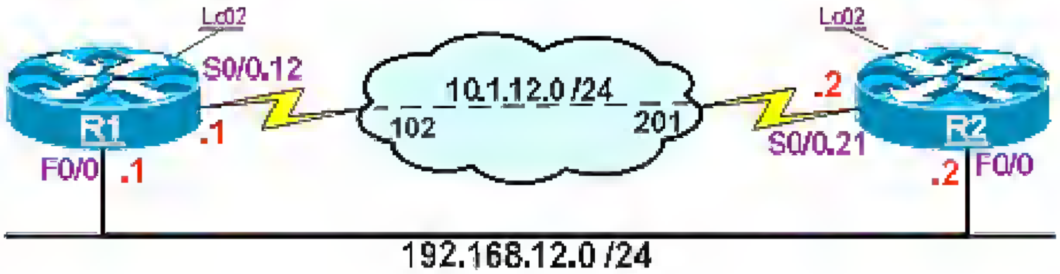

Lab Setup:

- Configure the frame-relay connection between R1 and R2 in a point to point manner.
- Configure the F0/0 interface of R1 and R2 in VLAN 12.
- Use the IP addressing chart for IP assignment.
- Configure RIPv2 to provide NLRI

IP addressing:

Router	Interface / IP address
R1	S0/0.12 =10.1.12.1 /24 Loopback0 = 1.1.1.1 /24 Loopback1 = 100.1.1.1 /24 F0/0 = 192.168.12.1 /24
R2	S0/0.21 = 10.1.12.2 /24 Loopback0 = 2.2.2.2 /24 Loopback1 = 200.2.2.2 /24 F0/0 = 192.168.12.2 /24

Task 1

R1 should be configured based on the following policy:

Traffic sourcing from 1.1.1.1 /24 should use the frame-relay connection, whereas, traffic sourcing from 100.1.1.1 /24 should take the F0/0 interface.

STEP 1

The following configuration identifies the source IP addresses of 1.1.1.1 and 100.1.1.1:

On R1

R1(config)#**Access-list 1 permit host 1.1.1.1**

R1(config)#**Access-list 2 permit host 100.1.1.1**

STEP 2

The following configuration, defines the actual policy using a route-map:

R1(config)#Route-map TST permit 10
R1(config-route-map)#Match ip address 1
R1(config-route-map)#**Set interface S0/0.12**

R1(config)#Route-map TST permit 20
R1(config-route-map)#Match ip address 2
R1(config-route-map)#**Set interface F0/0**

R1(config)#Route-map TST permit 30

STEP 3

In this step policy routing is enabled on the router, the following command enables the router to policy route packets that are sourced by the local router:

R1(config)#**Ip local policy route-map** TST

To test the configuration:

On R1

R1#Debug ip policy

R1#Ping 2.2.2.2 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms

IP: s=1.1.1.1 (local), d=2.2.2.2, len 100, **policy match**
IP: **route map TST, item 10, permit**
IP: **s=1.1.1.1 (local), d=2.2.2.2** (Serial0/0.12), len 100, **policy routed**
(The rest of the output is omitted)

R1#Ping 2.2.2.2 source 100.1.1.1

Route-map item 10
Route-map item 20

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 100.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

IP: s=100.1.1.1 (local), d=2.2.2.2, len 100, **policy match**
IP: **route map TST, item 20, permit**
IP: **s=100.1.1.1 (local), d=2.2.2.2** (FastEthernet0/0), len 100, **policy routed**
IP: **local to FastEthernet0/0** 192.168.12.2
IP: **s=100.1.1.1 (local), d=2.2.2.2,** len 100, **policy match**
IP: **route map TST, item 20, permit**

Note so far item 10 and 20 in the route-map has been matched in the output of the debug, the following Ping will match item 30 of this route-map:

R1#Ping 2.2.2.2 source F0/0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 192.168.12.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/8 ms

IP: s=192.168.12.1 (local), d=2.2.2.2, len 100, **policy match**
IP: route map TST, item 30, permit
IP: s=192.168.12.1 (local), d=2.2.2.2, len 100, policy rejected -- normal forwarding

Note the actual policy was rejected because it was routed using the routing table.

Task 2

Re-configure R1 based on the following policy:

- The communication from host 1.1.1.1 and 2.2.2.2 should traverse through the frame-relay connection.
- The communication from host 1.1.1.1 and 200.2.2.2 should traverse through the F0/0 interface.
- The communication from host 100.1.1.1 and 200.2.2.2 should traverse through the frame-relay connection.
- The communication from host 100.1.1.1 and 2.2.2.2 should traverse through the F0/0 interface.

Enter the following commands to remove the configuration from the previous step:

```
R1(config)#NO access-list 1
R1(config)#NO access-list 2
R1(config)#NO route-map TST
```

DO NOT remove the "Ip local policy route-map TST"

To configure the new access-lists:

```
R1(config)#ip access-list extended H1-2
R1(config-ext-nacl)#permit ip host 1.1.1.1 host 2.2.2.2

R1(config)#ip access-list extended H1-200
R1(config-ext-nacl)#permit ip host 1.1.1.1 host 200.2.2.2

R1(config)#ip access-list extended H100-200
R1(config-ext-nacl)#permit ip host 100.1.1.1 host 200.2.2.2

R1(config)#ip access-list extended H100-2
R1(config-ext-nacl)#permit ip host 100.1.1.1 host 2.2.2.2
```

To configure the new route-map:

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match ip addr H1-2
R1(config-route-map)#set interface S0/0.12

R1(config)#route-map TST permit 20
R1(config-route-map)#match ip addr H1-200
R1(config-route-map)#set interface F0/0
```

R1(config)#route-map TST permit 30
R1(config-route-map)#match ip addr H100-200
R1(config-route-map)#set interface S0/0.12

R1(config)#route-map TST permit 40
R1(config-route-map)#match ip addr H100-2
R1(config-route-map)#set interface F0/0

R1(config)#route-map TST permit 50

To test the configuration:

On R1

R1#Ping 2.2.2.2 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

IP: s=1.1.1.1 (local), d=2.2.2.2, len 100, policy match
IP: route map TST, item 10, permit
IP: s=1.1.1.1 (local), d=2.2.2.2 (Serial0/0.12), len 100, policy routed
IP: local to Serial0/0.12 10.1.12.2

R1#Ping 200.2.2.2 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms

IP: s=1.1.1.1 (local), d=200.2.2.2, len 100, policy match
IP: route map TST, item 20, permit
IP: s=1.1.1.1 (local), d=200.2.2.2 (FastEthernet0/0), len 100, policy routed
IP: local to FastEthernet0/0 192.168.12.2

R1#Ping 200.2.2.2 source 100.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.2.2.2, timeout is 2 seconds:

```
Packet sent with a source address of 100.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

IP: s=100.1.1.1 (local), d=200.2.2.2, len 100, policy match
IP: route map TST, item 30, permit
IP: s=100.1.1.1 (local), d=200.2.2.2 (Serial0/0.12), len 100, policy routed
IP: local to Serial0/0.12 10.1.12.2

R1#Ping 2.2.2.2 source 100.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 100.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

IP: s=100.1.1.1 (local), d=2.2.2.2, len 100, policy match
IP: route map TST, item 40, permit
IP: s=100.1.1.1 (local), d=2.2.2.2 (FastEthernet0/0), len 100, policy routed
IP: local to FastEthernet0/0 192.168.12.2

R1#Ping 2.2.2.2 source F0/0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 192.168.12.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

IP: s=192.168.12.1 (local), d=2.2.2.2, len 100, policy match
IP: route map TST, item 50, permit
IP: s=192.168.12.1 (local), d=2.2.2.2, len 100, policy rejected -- normal forwarding
```

Task 3

Remove all access-lists and the route-map from R1 before proceeding to the next task.

On R1

```
R1(config)#NO ip access-list extended H1-2
```

```
R1(config)#NO ip access-list extended H1-200
R1(config)#NO ip access-list extended H100-200
R1(config)#NO ip access-list extended H100-2

R1(config)#NO route-map TST
```

To verify the configuration:

On R1

R1#Show access-list

R1#Show route-map

R1#Show ip route rip

```
     2.0.0.0/24 is subnetted, 1 subnets
R       2.2.2.0 [120/1] via 192.168.12.2, 00:00:11, FastEthernet0/0
                [120/1] via 10.1.12.2, 00:00:21, Serial0/0.12
R    200.2.2.0/24 [120/1] via 192.168.12.2, 00:00:11, FastEthernet0/0
                  [120/1] via 10.1.12.2, 00:00:21, Serial0/0.12
```

On R2

R2#Show ip route rip

```
     1.0.0.0/24 is subnetted, 1 subnets
R       1.1.1.0 [120/1] via 192.168.12.1, 00:00:18, FastEthernet0/0
                [120/1] via 10.1.12.1, 00:00:02, Serial0/0.21
     100.0.0.0/24 is subnetted, 1 subnets
R       100.1.1.0 [120/1] via 192.168.12.1, 00:00:18, FastEthernet0/0
                  [120/1] via 10.1.12.1, 00:00:02, Serial0/0.21
```

Task 4

Configure R1 based on the following policy:

- If the size of the packet/s is up to 250 Bytes, they should traverse through the Frame-relay cloud.
- If the size of the packet/s is between 251 – 1500 Bytes, they should traverse through the F0/0 interface.

On R1

R1(config)#Route-map TST per 10
R1(config-route-map)#Match length 0 250
R1(config-route-map)#Set ip next-hop 10.1.12.2

R1(config)#Route-map TST per 20
R1(config-route-map)#Match length 251 1500
R1(config-route-map)#Set ip next-hop 192.168.12.2

R1(config)#Route-map TST per 30

To test the configuration:

On R1

R1#Ping 2.2.2.2 source 1.1.1.1 size 249

Type escape sequence to abort.
Sending 5, 249-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 76/76/77 ms

IP: s=1.1.1.1 (local), d=2.2.2.2, len 249, policy match
IP: route map TST, item 10, permit
IP: s=1.1.1.1 (local), d=2.2.2.2 (Serial0/0.12), len 249, policy routed
IP: local to Serial0/0.12 10.1.12.2

R1#Ping 2.2.2.2 source 1.1.1.1 size 251

Type escape sequence to abort.
Sending 5, 251-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/64/65 ms

IP: s=1.1.1.1 (local), d=2.2.2.2, len 251, policy match
IP: route map TST, item 20, permit
IP: s=1.1.1.1 (local), d=2.2.2.2 (FastEthernet0/0), len 251, policy routed

IP: local to FastEthernet0/0 192.168.12.2

```
R1#Ping 2.2.2.2 source 1.1.1.1 size 1501

Type escape sequence to abort.
Sending 5, 1501-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 344/346/349 ms

IP: s=1.1.1.1 (local), d=2.2.2.2, len 1501, policy match
IP: route map TST, item 30, permit
IP: s=1.1.1.1 (local), d=2.2.2.2, len 1501, policy rejected -- normal forwarding
IP: s=1.1.1.1 (local), d=2.2.2.2, len 1501, policy match
IP: route map TST, item 30, permit
```

Task 5

Remove the route-map from the previous step before proceeding to the next task.

On R1

```
R1(config)#NO route-map TST
```

To verify the configuration:

On R1

```
R1#Show route-map

R1#
```

Task 6

Re-configure R1 based on the following policy:

- All packets from any IP address destined to port 80 (HTTP) should traverse through the frame-relay cloud.
- All packets from any IP address destined to port 23 (Telnet) should traverse through port F0/0.

On R1

To configure the access-lists for identifying the traffic:

```
R1(config)#ip access-list extended HTTP
R1(config-ext-nacl)#permit tcp any any eq 80

R1(config)#ip access-list extended TELNET
R1(config-ext-nacl)#permit tcp any any eq 23
```

To configure the route-map for the policy:

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match ip addr HTTP
R1(config-route-map)#set interface S0/0.12

R1(config)#route-map TST permit 20
R1(config-route-map)#match ip addr TELNET
R1(config-route-map)#set interface F0/0

R1(config)#route-map TST permit 30
```

To test the configuration

On R1

```
R1#Telnet 200.2.2.2 80

Trying 200.2.2.2, 80 ... Open

IP: s=192.168.12.1 (local), d=200.2.2.2, len 44, policy match
IP: route map TST, item 10, permit
IP: s=192.168.12.1 (local), d=200.2.2.2 (Serial0/0.12), len 44, policy routed
IP: local to Serial0/0.12 10.1.12.2

R1#telnet 2.2.2.2

Trying 2.2.2.2 ... Open

Password required, but none set

IP: s=192.168.12.1 (local), d=2.2.2.2, len 44, policy match
IP: route map TST, item 20, permit
IP: s=192.168.12.1 (local), d=2.2.2.2 (FastEthernet0/0), len 44, policy routed
```

```
IP: local to FastEthernet0/0 192.168.12.2

R1#Telnet 2.2.2.2 8000

Trying 2.2.2.2, 8000 ...
% Connection refused by remote host

IP: s=192.168.12.1 (local), d=2.2.2.2, len 44, policy match
IP: route map TST, item 30, permit
IP: s=192.168.12.1 (local), d=2.2.2.2, len 44, policy rejected -- normal forwarding
```

Task 7

Erase the startup configuration and reload the routers before proceeding to the next lab.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

REDISTRIBUTION

Lab 1 – Basics of Redistribution-I

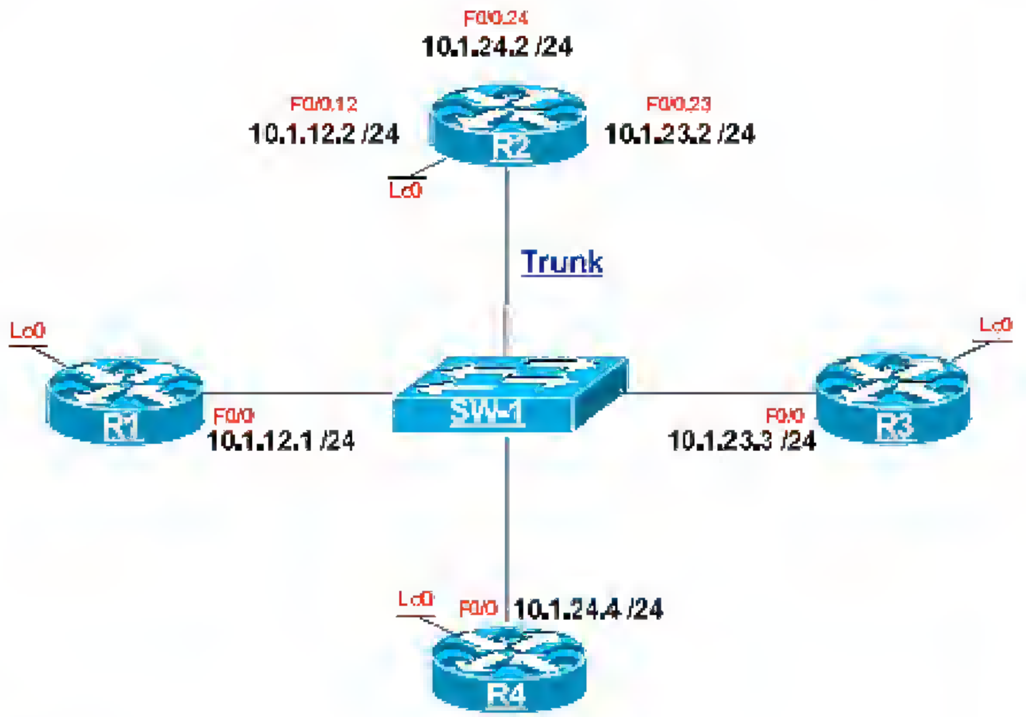

Lab Setup:

- Configure the F0/0 interface of R2 as a trunk providing routing for VLANs 12, 23 and 24.
- R1's F0/0 interface should be configured in VLAN 12.
- R3's F0/0 interface should be configured in VLAN 23.
- R4's F0/0 interface should be configured in VLAN 24.
- Use the IP addressing chart below for IP assignment.

IP addressing:

Router	Interface / IP address	Connecting to:
R2	F0/0.12 = 10.1.12.2 /24 F0/0.23 = 10.1.23.2 /24 F0/0.24 = 10.1.24.2 /24 Lo0 = 10.1.1.2 /32	R1's F0/0 R3's F0/0 R4's F0/0
R1	F0/0 = 10.1.12.1 /24 Loopback0 = 10.1.1.1 /32	R2's F0/0.12
R3	F0/0 = 10.1.23.3 /24 Loopback0 = 10.1.1.3 /32	R2's F0/0.23
R4	F0/0 = 10.1.24.4 /24 Loopback0 = 10.1.1.4 /32	R2's F0/0.24

Task 1

Configure the link between R2 and R3 to be in OSPF area 0. Configure Lo0 interface of R3 in area 2. Do not use the network command to accomplish this task.

On R2

```
R2(config)#int f0/0.23
R2(config-subif)#ip ospf 1 area 0
```

On R3:

```
R3(config)#int f0/0
R3(config-if)#ip ospf 1 area 0

R3(config-if)#int lo0
R3(config-if)#ip ospf 1 area 2
```

Note when running OSPF on a given interface by using the interface configuration mode command, the IOS starts the OSPF process automatically:

```
R3#Sh run | S router ospf 1

router ospf 1
 log-adjacency-changes
```

To verify the configuration:

R2 should see R3's Lo0 as "O IA" route in its routing table

On R2

R2#show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA    10.1.1.3/32 [110/2] via 10.1.23.3, 00:00:37, FastEthernet0/0.23
```

Task 2

Add the following Loopback interfaces to R3:

- Lo31: 192.168.31.3/24
- Lo32: 192.168.32.3/24
- Lo33: 192.168.33.3/24

Redistribute Lo31 and Lo33. Do not redistribute Lo32 into OSPF. Do not use "access-list" or "prefix-list" to accomplish the task.

The first step in resolving this task is to configure the loopback interfaces on R3, as follows:

On R3

```
R3(config)#int lo 31
R3(config-if)#ip addr 192.168.31.3 255.255.255.0

R3(config-if)#int lo 32
R3(config-if)#ip addr 192.168.32.3 255.255.255.0

R3(config-if)#int lo 33
R3(config-if)#ip addr 192.168.33.3 255.255.255.0
```

Since using "access-lists" or "prefix-lists" is not allowed, a route-map is configured and the required interfaces are matched using the "match interface" option in the "route-map" configuration mode.
Note the task states that Lo32 should NOT be redistributed, therefore, the "route-map" could be configured to deny Lo32 and permit the rest of the networks, or the "route-map" could simply permit the Lo31 and Lo33 interfaces and deny Lo32 from being redistributed.

Once the "route-map" is configured, its referenced in the "redistribute connected" command in the router configuration mode, as follows:

On R3

Option 1:

Note in the following option ONLY "Lo32" is denied, if "route-map TST permit 90" is not configured, the rest of the interfaces will also be denied from being redistributed, because there is an invisible implicit deny all statement at the end of every "route-map", therefore, the "route-map TST permit 90" is configured to permit the rest of the interfaces.

```
R3(config)#route-map TST deny 10
R3(config-route-map)#match interface Lo32
R3(config)#route-map TST permit 90
```

Option 2:

Note the following option is more specific and it should be used in this configuration. This option is more specific because it will not redistribute future directly connected routes in this process.

```
R3(config)#route-map TST permit 10
R3(config-route-map)#match interface Lo31 Lo33
```

In the final step the connected interfaces that are referenced in the "route-map" are redistributed, as follows:

```
R3(config-route-map)#router ospf 1
R3(config-router)#redistribute connected subnets route-map TST
```

The "subnets" keyword is required or else ONLY the classful networks are redistributed.

To verify the configuration:

On R3

R3#Show ip ospf da external | inc **Link State ID**

```
  Link State ID: 192.168.31.0 (External Network Number )
  Link State ID: 192.168.33.0 (External Network Number )
```

On R2

R2#Show ip route ospf | Inc E2

O E2 192.168.31.0/24 [110/20] via 10.1.23.3, 00:07:34, FastEthernet0/0.23
O E2 192.168.33.0/24 [110/20] via 10.1.23.3, 00:07:34, FastEthernet0/0.23

Task 3

Configure RIPv2 between R2 and R1.
On R2, redistribute OSPF routes to RIP.
R1 should not use the route to R3's Lo33 in its routing table. Do not use "distribute-list" or "offset-list" to accomplish the task; R2 is not allowed to filter any redistributed route from OSPF to RIP.

The first step is configure RIPv2 on both R1 and R2:

On R2

R2(config)#router rip
R2(config-router)#version 2
R2(config-router)#no auto-summary
R2(config-router)#net 10.0.0.0
R2(config-router)#Passive-interface F0/0.23
R2(config-router)#Passive-interface F0/0.24

Note in the above configuration the "Passive-interface" commands are required because the task requires RIPv2 to run between R1 and R2; the "Passive-interface" commands turn the specified interfaces into receive-only mode.

On R1:

R1(config)#router rip
R1(config-router)#version 2
R1(config-router)#no auto-summary
R1(config-router)#net 10.0.0.0

To verify the configuration:

On R2

R2#Show ip route rip

10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R 10.1.1.1/32 [120/1] via 10.1.12.1, 00:00:05, FastEthernet0/0.21

In the second step OSPF is redistributed into RIP on R2:

On R2

R2(config)#router rip
R2(config-router)#**redistribute** ospf 1

To verify the configuration:

On R1

R1#sh ip route rip

10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R 10.1.1.2/32 [120/1] via 10.1.12.2, 00:00:01, FastEthernet0/0
R 10.1.24.0/24 [120/1] via 10.1.12.2, 00:00:01, FastEthernet0/0
R 10.1.23.0/24 [120/1] via 10.1.12.2, 00:00:01, FastEthernet0/0

Note Lo0, Lo31 and Lo33 of R3 are NOT in the routing table of R1; only R2's interfaces are redistributed.

One of the biggest problems of redistribution is that each routing protocol has its own metric:

- RIP : hop count
- EIGRP : composite of bandwidth, delay, reliability, load and MTU
- OSPF : Cost which is based on bandwidth

When redistributing from OSPF into RIP routing protocol, what should be the metric? Well….there are many choices, and some of them are as follows:

- The metric can be configured such that it applies to all existing and future redistributed routes.
- The metric can be set separately on each configured redistribute command.
- The metric can be set based on usage of a "route-map".

One notable exception is directly connected routes, which RIP applies a default metric of 0.

To correct the problem, a metric of 3 is assigned to the OSPF routes redistributed into RIPv2, as follows:

On R2

```
R2(config)#router rip
R2(config-router)#redistribute ospf 1 metric 3
```

To verify the configuration:

On R1

R1#Sh ip route rip

```
R    192.168.31.0/24 [120/3] via 10.1.12.2, 00:00:11, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R       10.1.1.2/32      [120/1] via 10.1.12.2, 00:00:11, FastEthernet0/0
R       10.1.1.3/32      [120/3] via 10.1.12.2, 00:00:11, FastEthernet0/0
R       10.1.24.0/24     [120/1] via 10.1.12.2, 00:00:11, FastEthernet0/0
R       10.1.23.0/24     [120/1] via 10.1.12.2, 00:00:11, FastEthernet0/0
R    192.168.33.0/24 [120/3] via 10.1.12.2, 00:00:11, FastEthernet0/0
```

The output of the above show command verifies that all OSPF routes have been redistributed into RIPv2 on R1; the last step of this task requires filtering of R3's Lo33 (192.168.33.0/24), remember that the use of "distribute-list" or "Offset-list" is not allowed, therefore, a "route-map" is used to set the metric of interface Lo33 to infinity; this will cause R1 to poison that route.

The steps required to configure this task are as follows:

- A "prefix-list" is configured to identify the IP address of R3's Lo33
- A "route-map" is configured, the "prefix-list" from the previous step is referenced, and a "set metric 16" is configured to set the metric to infinity
- The redistribution is reconfigured to reference the "route map"

On R2

The first Step:

```
R2(config)#ip prefix-list Lo33 permit 192.168.33.0/24
```

The second step:

```
R2(config)#route-map TST 10
R2(config-route-map)#match ip address prefix-list Lo33
R2(config-route-map)#set metric 16
R2(config-route-map)#route-map TST 20
```

The third and the final step:

R2(config-route-map)#router rip
R2(config-router)#redistribute ospf 1 metric 3 route-map TST

To verify the configuration:

On R1

R1#Show ip route | Inc 192.168.33.0

Note the Lo33 of R3 is no longer in the routing table

R1#Show ip route rip

```
R    192.168.31.0/24 [120/3] via 10.1.12.2, 00:00:18, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R       10.1.1.2/32  [120/1] via 10.1.12.2, 00:00:18, FastEthernet0/0
R       10.1.1.3/32  [120/3] via 10.1.12.2, 00:00:18, FastEthernet0/0
R       10.1.24.0/24 [120/1] via 10.1.12.2, 00:00:18, FastEthernet0/0
R       10.1.23.0/24 [120/1] via 10.1.12.2, 00:00:18, FastEthernet0/0
```

R1#debug ip rip

```
RIP: received v2 update from 10.1.12.2 on FastEthernet0/0
      10.1.1.2/32 via 0.0.0.0 in 1 hops
      10.1.1.3/32 via 0.0.0.0 in 3 hops
      10.1.23.0/24 via 0.0.0.0 in 1 hops
      10.1.24.0/24 via 0.0.0.0 in 1 hops
      192.168.31.0/24 via 0.0.0.0 in 3 hops
      192.168.33.0/24 via 0.0.0.0 in 16 hops (inaccessible)
```

Note R1 receives R3's Lo33 with a metric of 16 and since RIP has a maximum hop count of 15, this route is inaccessible.

Task 4

Add the following loopback interfaces to R1 and R3:

- R1: Lo13: 192.168.13.1/24
- R3: Lo13: 192.168.13.3/24

Advertise R1's Lo13 into RIPv2 routing protocol; R3's Lo13 should be configured in OSPF area 2 and advertised with its correct mask.

Configure EIGRP adjacency between R2 and R4, these routers should be configured in Eigrp AS 100. Ensure that R2's routing table resembles the following; the composite metric can be any value:

R2#Show ip route eigrp

10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D EX 10.1.1.4/32
[170/156160] via 10.1.24.4, 00:00:52, FastEthernet0/0.24

Redistribute RIP into EIGRP. Ensure that R4 installs a route for network 192.168.13.0/24 into its routing table.

The first step:

Configuring a loopback 13 interface on R1 and another on R3:

On R1

R1(config)#int lo 13
R1(config-if)#ip addr 192.168.13.1 255.255.255.0

On R3:

R3(config)#int lo 13
R3(config-if)#ip addr 192.168.13.3 255.255.255.0

Second step:

Adding loopback 13 to OSPF area 2:

On R3

R3(config-router)#int lo13
R3(config-if)#ip ospf 1 area 2

On R3

R3#Show ip ospf interface loopback13

Loopback13 is up, line protocol is up

Internet Address 192.168.13.3/24, **Area 2**
Process ID 1, Router ID 10.1.1.3, **Network Type LOOPBACK**, Cost: 1
Enabled by interface config, including secondary ip addresses
Loopback interface is treated as a stub Host

Note when configuring loopback interfaces in OSPF, OSPF recognizes that this interface is a loopback interface, therefore, it treats it as a stub host.
As a stub host, OSPF and CPU operation/utilization is minimized because OSPF will NOT send any OSPF packets out of this interface or receive any OSPF packets.
OSPF advertises this interface as a host route with a mask of "/32", this behavior can be changed by changing the network type, as follows:

R3(config)#int lo13
R3(config-if)#**ip ospf network point-to-point**

To verify the configuration:

On R2

R2#Show ip route ospf | Inc 192.168.13.0

192.168.13.0/24 is subnetted, 1 subnets
O IA 192.168.13.0 [110/2] via 10.1.23.3, 00:01:59, FastEthernet0/0.23

Next step is to advertise Lo13 in RIPv2 on R1:

On R1

R1(config)#router rip
R1(config-router)#network 192.168.13.0

Note RIPv2 also recognizes that this interface is a Loopback interface the output of a "debug ip rip" will reveal that; RIP can be configured to minimize operation by adding the "Passive-interface Lo13" to the RIP configuration.

On R1

R1(config)#router rip
R1(config-router)#passive-interface loopback 13

To verify the configuration:

On R2

R2#Show ip route | Inc 192.168.13.0

```
     192.168.13.0/24 is subnetted, 1 subnets
O IA    192.168.13.0 [110/2] via 10.1.23.3, 00:12:48, FastEthernet0/0.23
```

Note the route does NOT appear as a RIP route; did RIP advertise the route? If so, what happened?

On R2

R2#Show ip rip database 192.168.13.0 255.255.255.0

```
192.168.13.0/24    redistributed
    [3] via 10.1.23.3, from 10.1.1.3,
```

Note the route is in RIP's database, but it is in the database as a "redistributed" route. It is in the database as "redistributed" because of redistribution of OSPF into RIP in task 3.

So where is the update from R1?
Maybe R1 does not send that update at all?

Note the output of the following debug reveals that the route is being advertised to R2:

On R2

R2#Debug ip rip

```
RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0.21 (10.1.12.2)
RIP: build update entries
    192.168.13.0/24 via 0.0.0.0, metric 3, tag 0
```

(The output is modified to show the advertisement for this route ONLY)

Note R2 received an update for network 192.168.13.0 /24 from R1, but it does not install it

into its RIP database, RIP database contains only redistributed routes and learned RIP routes which are actually installed in the router's routing table; since the "administrative distance" of RIP is 120 and OSPF's "Administrative distance" is 110, IOS prefers the OSPF route and rejects the RIP route.

To verify:

R3's Lo13 is shutdown and once again the R2's RIP's database is checked, as follows:

On R3

R3(config)#interface Loopback 13
R3(config-if)#Shut

On R2

R2#Show ip route 192.168.13.0

Routing entry for 192.168.13.0/24
 Known via "rip", distance 120, metric 1
 Redistributing via rip
 Last update from 10.1.12.1 on FastEthernet0/0.21, 00:00:28 ago
 Routing Descriptor Blocks:
 * 10.1.12.1, from 10.1.12.1, 00:00:28 ago, via FastEthernet0/0.21
 Route metric is 1, traffic share count is 1

R2#Show ip rip database 192.168.13.0 255.255.255.0

192.168.13.0/24
 [1] via 10.1.12.1, 00:00:15, FastEthernet0/0.21

Note the output of the above show commands reveal that network 192.168.13.0 /24 is a RIP route.

To verify and reveal the comparison of administrative distance of RIP versus OSPF by the IOS, the Lo13 interface of R3 is configured as "no shut" while "debug ip routing" is enabled on R2, as follows:

On R2

R2#debug ip routing

IP routing debugging is on

On R3

R3(config)#interface Loopback 13
R3(config-if)#No shut

On R2

RT: closer admin distance for 192.168.13.0, flushing 1 routes
RT: NET-RED 192.168.13.0/24
RT: SET_LAST_RDB for 192.168.13.0/24
 NEW rdb: via 10.1.23.3

RT: add 192.168.13.0/24 via 10.1.23.3, ospf metric [110/2]
RT: NET-RED 192.168.13.0/24

Note the first line of the above output reveals the comparison of the administrative distance of RIP and OSPF, and the "add 192.168.13.0/24 via 10.1.23.3, ospf metric [110/2] is what is injected in the routing table of this router, to verify that information:

On R2

R2#Sh ip route ospf | inc 192.168.13.0

O IA 192.168.13.0/24 [110/2] via 10.1.23.3, 00:08:30, FastEthernet0/0.23

In the next step EIGRP adjacency is created in AS 100 and the Loopback 0 interface of R4 is redistributed to resemble the output shown in this task:

On R2

R2(config)#router eigrp 100
R2(config-router)#no au
R2(config-router)#network 10.1.24.2 0.0.0.0

On R4

R4(config-route-map)#router eigrp 100
R4(config-router)#no au
R4(config-router)#netw 10.1.24.4 0.0.0.0

To see the loopback0 of R4 as an external Eigrp route, it must be redistributed into Eigrp 100, as follows:

The first step:

A route-map is created and Lo0 interface is matched:

R4(config)#route-map TST permit 10
R4(config-route-map)#match interface lo0

The second step:

The "Redistribute connected" command is configured referencing the "route-map", when redistributing routes into Eigrp, the metric should also be configured in the following order: Bandwidth, Delay, Relay, Load and MTU

The only exception is when redistributing connected and/or static routes.

Note EIGRP does not use the "MTU" as a parameter for metric calculation, but still a required value that must be configured.

R4(config)#router eigrp 100
R4(config-router)#redistribute connected route-map TST

To verify the configuration:

On R2

R2#Show ip route eigrp 100

```
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D EX    10.1.1.4/32 [170/156160] via 10.1.24.4, 00:07:28, FastEthernet0/0.24
```

****

Note when running Eigrp on an interface you should include the entire IP address of the given interface with an inverse mask option, or else you could run into problems/issues. In the beginning of this task the network command of R4 was configured in the most specific manner possible, if the network command was configured as "network 10.0.0.0", then, all interfaces that are configured in this major network would have been advertised as internal Eigrp route, and as a result of that the loopback 0 interface of R4 would have been injected as an internal and not external, the routing table of R2 would have resembled the following:

R2#sh ip route eigrp

```
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D       10.1.1.4/32 [90/156160] via 10.1.24.4, 00:00:31, FastEthernet0/0.24
```

****

The final step of this task:

RIPv2 is redistributed into Eigrp:

On R2

R2(config)#Router eigrp 100
R2(config-router)#redistribute rip metric 1 1 1 1 1

To verify the configuration:

On R4

R4#Show ip route eigrp

```
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
D EX    10.1.12.0/24 [170/2560002816] via 10.1.24.2, 00:01:23, FastEthernet0/0
D EX    10.1.1.2/32 [170/2560002816] via 10.1.24.2, 00:01:23, FastEthernet0/0
D EX    10.1.1.1/32 [170/2560002816] via 10.1.24.2, 00:01:23, FastEthernet0/0
D EX    10.1.23.0/24 [170/2560002816] via 10.1.24.2, 00:01:23, FastEthernet0/0
```

Note network 192.168.13.0 /24 has NOT been redistributed.

R4#Show ip route 192.168.13.0

% Network not in table

****************************** Important to remember ******************************

IOS will redistribute the 192.168.13.0 /24 network only if this network is in the routing table of R2 as a RIP route. Remember when redistributing routes into another routing protocol, let's say the routes from routing protocol A are to be redistributed into routing protocol B, the IOS will ONLY redistribute the routes that are in its routing table as A into routing protocol B's routing protocol.

Redistribution is NOT transitive, meaning that, when the routes from protocol A are redistributed into protocol B, the same routes can NOT be redistributed into routing protocol C.

****

To verify the routing table of R2 for network 192.168.13.0 /24:

On R2

R2#Show ip rip database 192.168.13.0 255.255.255.0

```
192.168.13.0/24    redistributed
    [3] via 10.1.23.3, from 10.1.1.3,
```

R2#Show ip route 192.168.13.0

```
Routing entry for 192.168.13.0/24
  Known via "ospf 1", distance 110, metric 2, type inter area
  Redistributing via rip
```

Advertised by rip metric 3 route-map TST
Last update from 10.1.23.3 on FastEthernet0/0.23, 01:18:05 ago
Routing Descriptor Blocks:
* 10.1.23.3, from 10.1.1.3, 01:18:05 ago, via FastEthernet0/0.23
  Route metric is 2, traffic share count is 1

The output of the above show command reveals that the route is present in RIP, but the route on R2 is known via OSPF, but we are redistributing RIP into EIGRP.

To accomplish this task, the administrative distance of 192.168.13.0 /24 is raised in OSPF to a number that is higher than OSPF's Administrative distance. As follows:

On R2 The desired AD value

Its important to note that this is the OSPF RID of R3

R2(config)#router ospf 1
R2(config-router)#distance 121 10.1.1.3 0.0.0.0

This is the inverse mask of R3's RID

To verify the configuration:

On R2:

Note the route is no longer in the routing table as an OSPF route:

R2#Sh ip route ospf

O E2 192.168.31.0/24 [121/20] via 10.1.23.3, 00:00:39, FastEthernet0/0.23
   10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA   10.1.1.3/32   [121/2] via 10.1.23.3, 00:00:39, FastEthernet0/0.23
O E2 192.168.33.0/24 [121/20] via 10.1.23.3, 00:00:39, FastEthernet0/0.23

Because of the adjusted AD, the route is now a RIP route:

R2#Show ip route rip

R   192.168.13.0/24 [120/1] via 10.1.12.1, 00:00:16, FastEthernet0/0.21
   10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
R   10.1.1.1/32   [120/1] via 10.1.12.1, 00:00:16, FastEthernet0/0.21

Note the output of the above show commands reveal that the configured AD is applied to all OSPF routes received from R3, even though the following show command reveals that network 192.168.13.0/24 is redistributed into R4's routing table; a better choice is a more specific approach.

To verify the configuration:

On R4

R4#Show ip route eigrp | Inc EX

```
D EX 192.168.13.0/24 [170/2560002816] via 10.1.24.2, 02:14:05, FastEthernet0/0
D EX    10.1.12.0/24 [170/2560002816] via 10.1.24.2, 02:22:54, FastEthernet0/0
D EX    10.1.1.2/32 [170/2560002816] via 10.1.24.2, 02:22:54, FastEthernet0/0
D EX    10.1.1.1/32 [170/2560002816] via 10.1.24.2, 02:22:54, FastEthernet0/0
D EX    10.1.23.0/24 [170/2560002816] via 10.1.24.2, 02:22:54, FastEthernet0/0
```

The following outlines the configuration of more specific approach on R2:

On R2

```
R2(config)#router ospf 1
R2(config-router)#no distance 121 10.1.1.3 0.0.0.0

R2(config-router)#distance 121 10.1.1.3 0.0.0.0 1

R2(config)#access-list 1 permit 192.168.13.0 0.0.0.255
```

To verify the configuration:

On R4

R4#Show ip route eigrp

```
D EX 192.168.13.0/24 [170/2560002816] via 10.1.24.2, 00:07:33, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
D EX    10.1.12.0/24    [170/2560002816] via 10.1.24.2, 02:46:04, FastEthernet0/0
D EX    10.1.1.2/32     [170/2560002816] via 10.1.24.2, 02:46:04, FastEthernet0/0
D EX    10.1.1.1/32     [170/2560002816] via 10.1.24.2, 00:07:33, FastEthernet0/0
D EX    10.1.23.0/24    [170/2560002816] via 10.1.24.2, 02:46:04, FastEthernet0/0
```

Task 5

Configure the following Loopback interface on R1:

Lo40: 172.16.0.1/24

R1 should be configured to advertise this loopback interface in RIPv2

Ensure that the **ONLY** 172.16.X.X route present in R2's routing table is 172.16.0.0/24

R3 should see the following route: 172.16.0.0/16.

Step one:

Configuring Loopback 40 on R1 and advertise this network in RIPv2:

On R1

R1(config)#int lo 40
R1(config-if)#ip addr 172.16.0.1 255.255.255.0

R1(config-if)#router rip
R1(config-router)#Netw 172.16.0.0

To verify the configuration:

On R2

R2#Sh ip route rip

```
R    192.168.13.0/24 [120/1] via 10.1.12.1, 00:00:13, FastEthernet0/0.21
     172.16.0.0/24 is subnetted, 1 subnets
R       172.16.0.0       [120/1] via 10.1.12.1, 00:00:13, FastEthernet0/0.21
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
R       10.1.1.1/32      [120/1] via 10.1.12.1, 00:00:13, FastEthernet0/0.21
```

Note R2 has the route based on the requirement of the task.
To perform the last step of this task, R2 should be configured to redistribute RIPv2 into OSPF, as follows:

On R2

Note by default OSPF assigns a metric of 20 and a metric-type of 2 to all redistributed routes, therefore, there is no need to assign the cost from OSPF's perspective when redistributing routes.

R2(config)#router ospf 1
R2(config-router)#**redistribute rip subnets**

To verify the configuration:

On R3

R3#Show ip route ospf

```
     172.16.0.0/24 is subnetted, 1 subnets
O E2    172.16.0.0 [110/20] via 10.1.23.2, 00:00:52, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E2    10.1.12.0/24 [110/20] via 10.1.23.2, 00:00:52, FastEthernet0/0
O E2    10.1.1.2/32   [110/20] via 10.1.23.2, 00:00:52, FastEthernet0/0
O E2    10.1.1.1/32   [110/20] via 10.1.23.2, 00:00:52, FastEthernet0/0
O E2    10.1.24.0/24 [110/20] via 10.1.23.2, 00:00:52, FastEthernet0/0
```

Note this task specified that network 172.16.0.0 should appear in the routing table of R3 as a "/16", to accomplish this task this network is summarized, as follows:

On R2

```
R2(config)#router ospf 1
R2(config-router)#summary-address 172.16.0.0 255.255.0.0
```

To verify the configuration:

On R3

R3#Sh ip route ospf

```
O E2 172.16.0.0/16  [110/20] via 10.1.23.2, 00:00:18, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E2    10.1.12.0/24 [110/20] via 10.1.23.2, 00:10:29, FastEthernet0/0
O E2    10.1.1.2/32   [110/20] via 10.1.23.2, 00:10:29, FastEthernet0/0
O E2    10.1.1.1/32   [110/20] via 10.1.23.2, 00:10:29, FastEthernet0/0
O E2    10.1.24.0/24 [110/20] via 10.1.23.2, 00:10:29, FastEthernet0/0
```

Note this fulfills the requirement of the last task, but remember this task specified that R2 should ONLY have 172.16.0.0 /24 in its routing table.

To verify the configuration:

On R2

R2#Show ip route | Inc 172.16

```
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

```
R       172.16.0.0/24 [120/1] via 10.1.12.1, 00:00:15, FastEthernet0/0.21
O       172.16.0.0/16 is a summary, 00:09:41, Null0
```

R2 has two 172.16.X.X routes in its routing table; 172.16.0.0/24 which is learned from RIP and 172.16.0.0/16 which is the direct result of summarization, remember that when we summarize in OSPF, the IOS will inject a discard route to avoid forwarding loops, the injected discard route can be for internal OSPF routes that were summarized or external OSPF routes that were summarized.

If the discard route is internal and it needs to be removed, then the "No discard-route internal" command can be used, but in this case its for external routes, therefore, it can be removed using the following configuration:

On R2

```
R2(config)#router ospf 1
R2(config-router)#no discard-route external
```

To verify the configuration:

On R2

R2#Show ip route | inc 172.16

```
     172.16.0.0/24 is subnetted, 1 subnets
R       172.16.0.0 [120/1] via 10.1.12.1, 00:00:27, FastEthernet0/0.21
```

On R3

R3#Sh ip route ospf

```
O E2 172.16.0.0/16 [110/20] via 10.1.23.2, 00:24:19, FastEthernet0/0
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E2    10.1.12.0/24 [110/20] via 10.1.23.2, 00:34:30, FastEthernet0/0
O E2    10.1.1.2/32   [110/20] via 10.1.23.2, 00:34:30, FastEthernet0/0
O E2    10.1.1.1/32   [110/20] via 10.1.23.2, 00:34:30, FastEthernet0/0
O E2    10.1.24.0/24 [110/20] via 10.1.23.2, 00:34:30, FastEthernet0/0
```

Task 6

Erase the startup configuration and reload the routers, you should also delete the "config.text" and "Vlan.dat" on SW1 and reload this switch before proceeding to the next lab.

Lab 2 – Basics of Redistribution-II

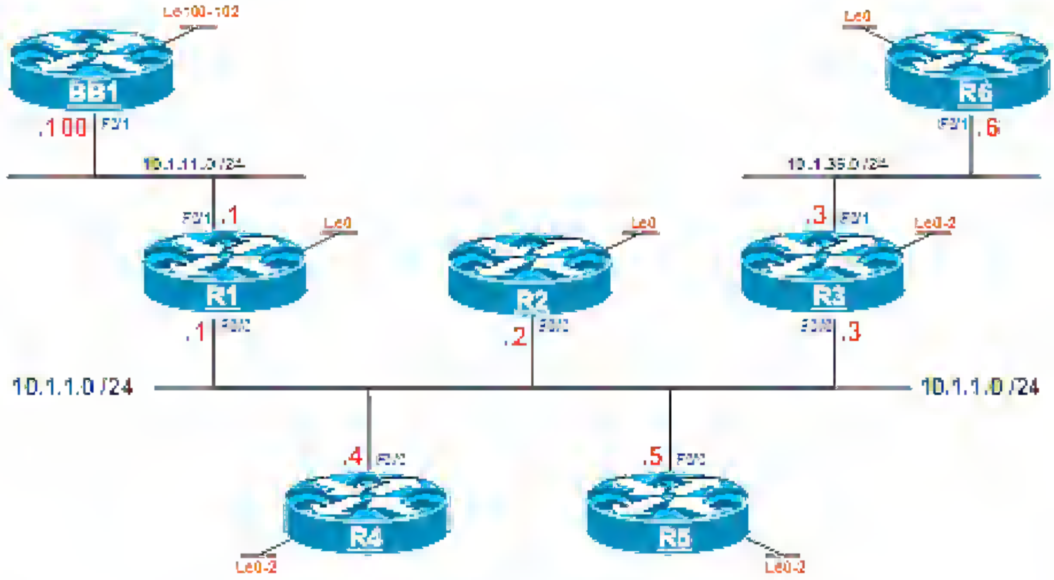

Lab Setup:

- The F0/0 interface of R1, R2, R3, R4 and R5 should be configured in VLAN 100
- The F0/1 interface of R1 and BB1 should be configured in VLAN 200
- The F0/1 interface of R3 and R6 should be configured in VLAN 300
- Use the following IP addressing chart for IP addressing assignment

IP addressing chart:

Router	Interface & IP Addressing:
R1	F0/0 – 10.1.1.1 /24 F0/1 – 10.1.11.1 /24 Lo0 – 1.1.1.1 /24
R2	F0/0 – 10.1.1.2 /24 Lo0 – 2.2.2.2 /24
R3	F0/0 – 10.1.1.3 /24 F0/1 – 10.1.36.3 /24 Lo0 – 3.3.3.3 /24 Lo1 – 30.3.3.3 /24 Lo2 – 33.3.3.3 /24
R4	F0/0 – 10.1.1.4 /24 Lo0 – 4.4.4.4 /24 Lo1 – 40.4.4.4 /24 Lo2 – 44.4.4.4 /24
R5	F0/0 – 10.1.1.5 /24 Lo0 – 5.5.5.5 /24 Lo1 – 50.5.5.5 /24 Lo2 – 55.5.5.5 /24
R6	F0/1 – 10.1.36.6 /24 Lo0 – 6.6.6.6 /24
BB1	F0/1 – 10.1.11.100 /24 Lo100 – 100.1.1.1 /24 Lo101 – 101.1.1.1 /24 Lo102 – 102.1.1.1 /24

Task 1

Configure the F0/0 interface of R3, R4, R5 and Lo0, 1 and Lo2 of R4 and R5 and Lo1 of R3 in OSPF area 0. Loopback interfaces should be advertised with their correct mask; assign the following OSPF costs to the loopback interfaces of R4 and R5:

Router	Interface & OSPF Cost
R4	Lo0 – Cost 10 Lo1 – Cost 20 Lo2 – Cost 30
R5	Lo0 – Cost 10 Lo1 – Cost 20 Lo2 – Cost 30

When configuring OSPF in this task, OSPF Process ID of 1 should be used

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 30.3.3.3 0.0.0.0 area 0
```

On R4

```
R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
R4(config-if)#ip ospf cost 10

R4(config-router)#int lo1
R4(config-if)#ip ospf network point-to-point
R4(config-if)#ip ospf cost 20

R4(config-router)#int lo2
R4(config-if)#ip ospf network point-to-point
R4(config-if)#ip ospf cost 30

R4(config)#router ospf 1
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 40.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 44.4.4.4 0.0.0.0 area 0
```

On R5

```
R5(config-router)#int lo0
R5(config-if)#ip ospf network point-to-point
R5(config-if)#ip ospf cost 10

R5(config-router)#int lo1
R5(config-if)#ip ospf network point-to-point
R5(config-if)#ip ospf cost 20

R5(config-router)#int lo2
R5(config-if)#ip ospf network point-to-point
R5(config-if)#ip ospf cost 30

R5(config)#router ospf 1
R5(config-router)#netw 10.1.1.5 0.0.0.0 area 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 0
R5(config-router)#netw 50.5.5.5 0.0.0.0 area 0
```

R5(config-router)#netw 55.5.5.5 0.0.0.0 area 0

To verify the configuration:

On R3

R3#Show ip route ospf | Inc O

```
O    50.5.5.0 [110/21] via 10.1.1.5, 00:03:47, FastEthernet0/0
O    4.4.4.0 [110/11] via 10.1.1.4, 00:03:47, FastEthernet0/0
O    55.5.5.0 [110/31] via 10.1.1.5, 00:03:47, FastEthernet0/0
O    5.5.5.0 [110/11] via 10.1.1.5, 00:03:47, FastEthernet0/0
O    40.4.4.0 [110/21] via 10.1.1.4, 00:03:47, FastEthernet0/0
O    44.4.4.0 [110/31] via 10.1.1.4, 00:03:47, FastEthernet0/0
```

Task 2

Configure Eigrp AS 100 on the F0/0 interface of R3, R2, R1; F0/1 interface of R1 and BB1; Lo0, Lo1 and Lo2 of BB1; Lo0 of R1, R2 and R3.

On R3

```
R3(config)#router eigrp 100
R3(config-router)#no au
R3(config-router)#netw 10.1.1.3 0.0.0.0
R3(config-router)#netw 3.3.3.3 0.0.0.0
```

On R2

```
R2(config)#router eigrp 100
R2(config-router)#no au
R2(config-router)#netw 10.1.1.2 0.0.0.0
R2(config-router)#netw 2.2.2.2 0.0.0.0
```

On R1

```
R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#netw 10.1.1.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

R1(config-router)#netw 10.1.11.1 0.0.0.0

On BB1

BB1(config)#router eigrp 100
BB1(config-router)#no au
BB1(config-router)#netw 10.1.11.100 0.0.0.0
BB1(config-router)#netw 100.1.1.1 0.0.0.0
BB1(config-router)#netw 101.1.1.1 0.0.0.0
BB1(config-router)#netw 102.1.1.1 0.0.0.0

To verify the configuration:

On R3

R3#Show ip route eigrp | Inc D

```
D    102.1.1.0 [90/158720] via 10.1.1.1, 00:04:23, FastEthernet0/0
D    1.1.1.0 [90/156160] via 10.1.1.1, 00:06:44, FastEthernet0/0
D    2.2.2.0 [90/156160] via 10.1.1.2, 00:00:07, FastEthernet0/0
D    100.1.1.0 [90/158720] via 10.1.1.1, 00:04:33, FastEthernet0/0
D    101.1.1.0 [90/158720] via 10.1.1.1, 00:04:29, FastEthernet0/0
D    10.1.11.0 [90/30720] via 10.1.1.1, 00:06:34, FastEthernet0/0
```

Task 3

Configure another OSPF routing domain using OSPF process ID of 36 on the F0/1 interface of R3 and R6; Lo0 interface of R6 and Lo2 of R3. These loopback interfaces should be advertised with their correct mask

On R3

R3(config-if)#int lo2
R3(config-if)#ip ospf net point-to-point

R3(config)#router ospf 36
R3(config-router)#netw 10.1.36.3 0.0.0.0 area 0
R3(config-router)#netw 33.3.3.3 0.0.0.0 area 0

On R6

R6(config)#int lo0
R6(config-if)#ip ospf netw point-to-point

R6(config)#router ospf 36
R6(config-router)#netw 10.1.36.6 0.0.0.0 area 0
R6(config-router)#netw 6.6.6.6 0.0.0.0 area 0

To verify the configuration:

On R3

R3#Show ip route ospf 36 | Inc O

O 6.6.6.0 [110/2] via 10.1.36.6, 00:00:59, FastEthernet0/1

Task 4

Configure R3 to redistribute Eigrp 100 into OSPF 1 such that networks 1.1.1.0 /24 and 2.2.2.0 /24 will have a tag of 111 and 222 respectively, the rest of the routes should have a route tag of 333.

On R3

R3(config)#access-list 1 permit 1.1.1.0 0.0.0.255
R3(config)#access-list 2 permit 2.2.2.0 0.0.0.255

R3(config)#route-map **TST** permit 10
R3(config-route-map)#match ip addr 1
R3(config-route-map)#**set tag 111**

R3(config)#route-map **TST** permit 20
R3(config-route-map)#match ip addr 2
R3(config-route-map)#**set tag 222**

R3(config)#route-map **TST** permit 30
R3(config-route-map)#**set tag 333**

R3(config)#router ospf 1
R3(config-router)#redistribute eigrp 100 route-map **TST** subnets

To verify the configuration:

On R3

R3#Show ip ospf 1 database

```
        OSPF Router with ID (33.3.3.3) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count

33.3.3.3        33.3.3.3        179         0x80000008 0x008002 2
44.4.4.4        44.4.4.4        641         0x80000005 0x00487A 4
55.5.5.5        55.5.5.5        467         0x80000005 0x00F791 4

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.3        33.3.3.3        665         0x80000003 0x007AD7

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum   Tag
1.1.1.0         33.3.3.3        179         0x80000001 0x00C531   111
2.2.2.0         33.3.3.3        179         0x80000001 0x0085FD   222
3.3.3.0         33.3.3.3        179         0x80000001 0x00928C   333
10.1.11.0       33.3.3.3        179         0x80000001 0x008C78   333
100.1.1.0       33.3.3.3        179         0x80000001 0x006450   333
101.1.1.0       33.3.3.3        179         0x80000001 0x00575C   333
102.1.1.0       33.3.3.3        192         0x80000001 0x004A68   333
```

On R4 or R5:

R4#Show ip route ospf | Inc O

```
R4#Show ip route ospf | Inc E2
O E2    102.1.1.0 [110/20] via 10.1.1.1, 00:03:07, FastEthernet0/0
O E2    1.1.1.0 [110/20] via 10.1.1.1, 00:03:07, FastEthernet0/0
O E2    2.2.2.0 [110/20] via 10.1.1.2, 00:03:07, FastEthernet0/0
O E2    100.1.1.0 [110/20] via 10.1.1.1, 00:03:07, FastEthernet0/0
O E2    3.3.3.0 [110/20] via 10.1.1.3, 00:03:07, FastEthernet0/0
O E2    101.1.1.0 [110/20] via 10.1.1.1, 00:03:07, FastEthernet0/0
O E2    10.1.11.0 [110/20] via 10.1.1.1, 00:03:07, FastEthernet0/0
```

On R4

```
R4#Show ip ospf database | B Type-5 AS External

            Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum    Tag
1.1.1.0         33.3.3.3        368         0x80000001 0x00C531    111
2.2.2.0         33.3.3.3        368         0x80000001 0x0085FD    222
3.3.3.0         33.3.3.3        368         0x80000001 0x00928C    333
10.1.11.0       33.3.3.3        368         0x80000001 0x008C78    333
100.1.1.0       33.3.3.3        368         0x80000001 0x006450    333
101.1.1.0       33.3.3.3        368         0x80000001 0x00575C    333
102.1.1.0       33.3.3.3        368         0x80000001 0x004A68    333
```

Task 5

Configure R4 to filter all routes that are tagged with 111, you should NOT use an access-list or a prefix-list to accomplish this task.

On R4

```
R4(config)#route-map TST deny 10
R4(config-route-map)#match tag 111

R4(config)#route-map TST permit 20

R4(config)#router ospf 1
R4(config-router)#distribute-list route-map TST in
```

To verify the configuration:

On R4

```
R4#Show ip route ospf | Inc E2

O E2    102.1.1.0 [110/20] via 10.1.1.1, 00:01:52, FastEthernet0/0
O E2    2.2.2.0 [110/20] via 10.1.1.2, 00:01:52, FastEthernet0/0
O E2    100.1.1.0 [110/20] via 10.1.1.1, 00:01:52, FastEthernet0/0
O E2    3.3.3.0 [110/20] via 10.1.1.3, 00:01:52, FastEthernet0/0
O E2    101.1.1.0 [110/20] via 10.1.1.1, 00:01:52, FastEthernet0/0
O E2    10.1.11.0 [110/20] via 10.1.1.1, 00:01:52, FastEthernet0/0
```

Note Network 1.1.1.0 /24 is no longer in the routing table of R1, but you should remember that it is still in the database of this router, the output of the following show command reveals that information:

R4#**Show ip ospf database external | Inc _1.1.1.0**

Link State ID: 1.1.1.0 (External Network Number)

Task 6

Configure R5 to filter all routes that are tagged with 222, you should NOT use an access-list or a prefix-list to accomplish this task.

On R5

```
R5(config)#route-map TST deny 10
R5(config-route-map)#match tag 222

R5(config)#route-map TST permit 20

R5(config)#router ospf 1
R5(config-router)#distribute-list route-map TST in
```

To verify the configuration:

On R5

R5#Show ip route ospf | Inc E2

```
O E2    102.1.1.0 [110/20] via 10.1.1.1, 00:01:55, FastEthernet0/0
O E2    1.1.1.0 [110/20] via 10.1.1.1, 00:01:55, FastEthernet0/0
O E2    100.1.1.0 [110/20] via 10.1.1.1, 00:01:55, FastEthernet0/0
O E2    3.3.3.0 [110/20] via 10.1.1.3, 00:01:55, FastEthernet0/0
O E2    101.1.1.0 [110/20] via 10.1.1.1, 00:01:55, FastEthernet0/0
O E2    10.1.11.0 [110/20] via 10.1.1.1, 00:01:55, FastEthernet0/0
```

Task 7

Configure R3 to redistribute OSPF 1 into Eigrp 100 such that network 4.4.4.0 /24 and 5.5.5.0 /24 are tagged with 444 and 555 respectively

On R3

R3(config)#access-list 4 permit 4.4.4.0 0.0.0.255
R3(config)#access-list 5 permit 5.5.5.0 0.0.0.255

R3(config)#route-map TST1 permit 10
R3(config-route-map)#match ip addr 4
R3(config-route-map)#set tag 444

R3(config)#route-map TST1 permit 20
R3(config-route-map)#match ip addr 5
R3(config-route-map)#set tag 555

R3(config)#route-map TST1 permit 30

R3(config)#router eigrp 100
R3(config-router)#redistribute ospf 1 route-map TST1 metric 1 1 1 1 1

To verify the configuration:

On R3

R3#Sh ip eigrp topology | Inc 444|555

P 4.4.4.0/24, 1 successors, FD is 2560000256, tag is 444
P 5.5.5.0/24, 1 successors, FD is 2560000256, tag is 555

Task 8

Configure R2 to filter network 4.4.4.0 /24, do not use an access-list or a prefix-list to accomplish this task.

On R2

R2#Sh ip eigrp topology | Inc 444

P 4.4.4.0/24, 1 successors, FD is 2560002816, tag is 444

Note R2 sees the tag that was assigned in the previous task.

R2(config)#route-map **TST deny 10**
R2(config-route-map)#**match tag 444**

R2(config)#route-map **TST permit 20**

R2(config)#router eigrp 100
R2(config-router)#distribute-list route-map TST in

To verify the configuration:

On R2

R2#Show ip route eigrp | Inc EX

```
D EX    50.5.5.0 [170/2560002816] via 10.1.1.3, 00:10:26, FastEthernet0/0
D EX    55.5.5.0 [170/2560002816] via 10.1.1.3, 00:10:26, FastEthernet0/0
D EX    5.5.5.0 [170/2560002816] via 10.1.1.3, 00:07:17, FastEthernet0/0
D EX    40.4.4.0 [170/2560002816] via 10.1.1.3, 00:10:26, FastEthernet0/0
D EX    44.4.4.0 [170/2560002816] via 10.1.1.3, 00:10:26, FastEthernet0/0
D EX    30.3.3.0 [170/2560002816] via 10.1.1.3, 00:10:26, FastEthernet0/0
```

Task 9

Configure R5 to filter all routes that originated by R4. DO NOT use a prefix-list or route tags to accomplish this task.

On R5

Note R5 receives the following three routes from R4:

R5#Show ip route | Inc 10.1.1.4

```
O       4.4.4.0 [110/11] via 10.1.1.4, 00:21:48, FastEthernet0/0
O       40.4.4.0 [110/21] via 10.1.1.4, 00:21:48, FastEthernet0/0
O       44.4.4.0 [110/31] via 10.1.1.4, 00:21:48, FastEthernet0/0
```

To configure the task:

The "match ip route-source" can be used to accomplish this task; IN OSPF THE ROUTE-SOURCE IS NOT THE NEXT HOP IP ADDRESS, IT IS THE ROUTER ID OF THE ROUTER THAT ADVERTISED THE ROUTES; to see the router id of R4:

R4#Sh ip ospf | Inc ID

Routing Process "ospf 1" with ID 44.4.4.4

R5(config)#access-list 4 permit 44.4.4.4

In task 6 the following route-map was created and applied inbound on R5:

R5(config)#route-map TST deny 10
R5(config-route-map)#match tag 222

R5(config)#route-map TST permit 20

This task will add the following line to the same route-map:

R5(config)#route-map TST1 deny 10
R5(config-route-map)#match ip route-source 4

As a result the new route-map should look like the following:

On R5

R5(config)#route-map TST deny 10
R5(config-route-map)#match tag 222

R5(config)#route-map TST1 deny 15
R5(config-route-map)#match ip route-source 4

R5(config)#route-map TST permit 20

Note the reason we have the sequence numbers in a route-map is so we can add configuration lines in the middle of a route-map, in this case route-map sequence number 15 is added between route-map 10 and 20.

Since the route-map is already applied, no further configuration is necessary

To verify the configuration:

On R5

R5#Show ip route | Inc 10.1.1.4

R5#

Note the networks that were advertised by R4 are filtered.

Task 10

Configure R3 to redistribute all routes from "OSPF 1" and all Eigrp 100 routing domain into "OSPF 36" routing domain using the following policy:

- The routes from "OSPF 1" routing domain with an OSPF cost of 11 should be redistributed as "E2" OSPF routes, these routes should have a tag of 11
- The routes from "OSPF 1" routing domain with an OSPF cost of 21 should be redistributed as "E2" OSPF routes, these routes should have a tag of 21
- The routes from "OSPF 1" routing domain with an OSPF cost of 31 should be redistributed as "E2" OSPF routes, these routes should have a tag of 31
- Eigrp 100 routes should be redistributed as "E1" routes with a route tag of 99.
- DO NOT use an access-list or prefix-list to accomplish this task

On R3

```
R3(config)#route-map TST9 permit 10
R3(config-route-map)#match source-protocol OSPF 1
R3(config-route-map)#match metric 11
R3(config-route-map)#set tag 11

R3(config)#route-map TST9 permit 20
R3(config-route-map)#match source-protocol OSPF 1
R3(config-route-map)#match metric 21
R3(config-route-map)#set tag 21

R3(config)#route-map TST9 permit 30
R3(config-route-map)#match source-protocol ospf 1
R3(config-route-map)#match metric 31
R3(config-route-map)#set tag 31
```

R3(config)#route-map TST9 permit 40

R3(config)#route-map **Eigrp** permit 10
R3(config-route-map)#**match source-protocol Eigrp 100**
R3(config-route-map)#**set tag 99**
R3(config-route-map)#**set metric-type type-1**

R3(config)#router ospf 36
R3(config-router)#redistribute ospf 1 route-map TST9 subnets
R3(config-router)#redistribute eigrp 100 subnets route-map **Eigrp**

To verify the configuration:

On R6

R6#Show ip route ospf | Inc E1

```
O E1    102.1.1.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    1.1.1.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    2.2.2.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    100.1.1.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    3.3.3.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    101.1.1.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    10.1.11.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
O E1    10.1.1.0 [110/21] via 10.1.36.3, 00:09:25, FastEthernet0/1
```

R6#Show ip route ospf | Inc E2

```
O E2    50.5.5.0 [110/21] via 10.1.36.3, 00:14:15, FastEthernet0/1
O E2    4.4.4.0 [110/11] via 10.1.36.3, 00:14:15, FastEthernet0/1
O E2    55.5.5.0 [110/31] via 10.1.36.3, 00:14:15, FastEthernet0/1
O E2    5.5.5.0 [110/11] via 10.1.36.3, 00:14:15, FastEthernet0/1
O E2    40.4.4.0 [110/21] via 10.1.36.3, 00:14:15, FastEthernet0/1
O E2    44.4.4.0 [110/31] via 10.1.36.3, 00:14:15, FastEthernet0/1
O E2    30.3.3.0 [110/1] via 10.1.36.3, 00:12:14, FastEthernet0/1
```

R6#Show ip ospf database | Inc 11_

```
4.4.4.0       30.3.3.3      999       0x80000001 0x0075F5 11
5.5.5.0       30.3.3.3      999       0x80000001 0x005117 11
```

R6#Show ip ospf database | Inc 21_

```
40.4.4.0      30.3.3.3      1069      0x80000001 0x00B87A 21
```

```
50.5.5.0        30.3.3.3        1069        0x80000001 0x001F08  21

R6#Show ip ospf database | Inc 31_

44.4.4.0        30.3.3.3        1093        0x80000001 0x009D7D  31
55.5.5.0        30.3.3.3        1093        0x80000001 0x00F617   31

R6#Show ip ospf database | Inc 99_

1.1.1.0          30.3.3.3        926         0x80000002 0x00ECA5   99
2.2.2.0          30.3.3.3        926         0x80000002 0x00C8C6    99
3.3.3.0          30.3.3.3        926         0x80000002 0x00A4E7    99
10.1.1.0         30.3.3.3        926         0x80000002 0x007712     99
10.1.11.0        30.3.3.3        926         0x80000002 0x000976     99
100.1.1.0        30.3.3.3        926         0x80000002 0x00E04E    99
101.1.1.0        30.3.3.3        926         0x80000002 0x00D35A   99
102.1.1.0        30.3.3.3        926         0x80000002 0x00C666    99
```

Task 11

Configure R3 such that Eigrp routes that have a composite metric of 156160 to 158720 are NOT redistributed into "OSPF 36" routing domain. None of the previous configurations should be removed or overridden to accomplish this task.

This task is asking for the routes that have a composite metric of 156160 to 158720, the formula to calculate the metric and the deviation value is as follows:

The two numbers are added and then divided by 2:
156160 + 158720 = 314880
314880 / 2 = 157440 → This gives us the metric value

To calculate the deviation value:
Subtract the start of the range number (The lower value or the "*from*" value, in this case 156160) from the "*to*" value, in this case 158720 and then divide the result by two:

158720 – 156160 = 2560
2560 / 2 = 1280
Therefore, the "match metric" command will have the following values:
Match metric 157440 +-1280

To test the values:

157440 - 1280 = 156160 Note this is the start of the range
157440 + 1280 = 158720 this is the top of the range.

On R3

R3(config)#route-map TASK-10 deny 10
R3(config-route-map)#match metric 157440 +- 1280

R3(config)#route-map TASK-10 permit 20

R3(config-route-map)#router eigrp 100
R3(config-router)#distribute-list route-map TASK-10 in

To verify the configuration:

On R3

R3#Show ip route eigrp

```
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.11.0 [90/30720] via 10.1.1.1, 02:15:28, FastEthernet0/0
```

Note if it is not in the routing table of R3 as an Eigrp route it will NOT get redistributed.

On R6

R6#Show ip ospf database | Inc 99_

```
3.3.3.0       30.3.3.3      1373      0x80000003 0x00A2E8  99
10.1.1.0      30.3.3.3      1373      0x80000003 0x007513   99
10.1.11.0     30.3.3.3      1373      0x80000003 0x000777   99
```

Note networks 3.3.3.0 and 10.1.1.0 are directly connected links to R3 and therefore, the composite metric of these two routes are NOT within the stated range.

Task 12

Erase the startup configuration and reload the routers before proceeding to the next task.

Lab 3 – Redistribution

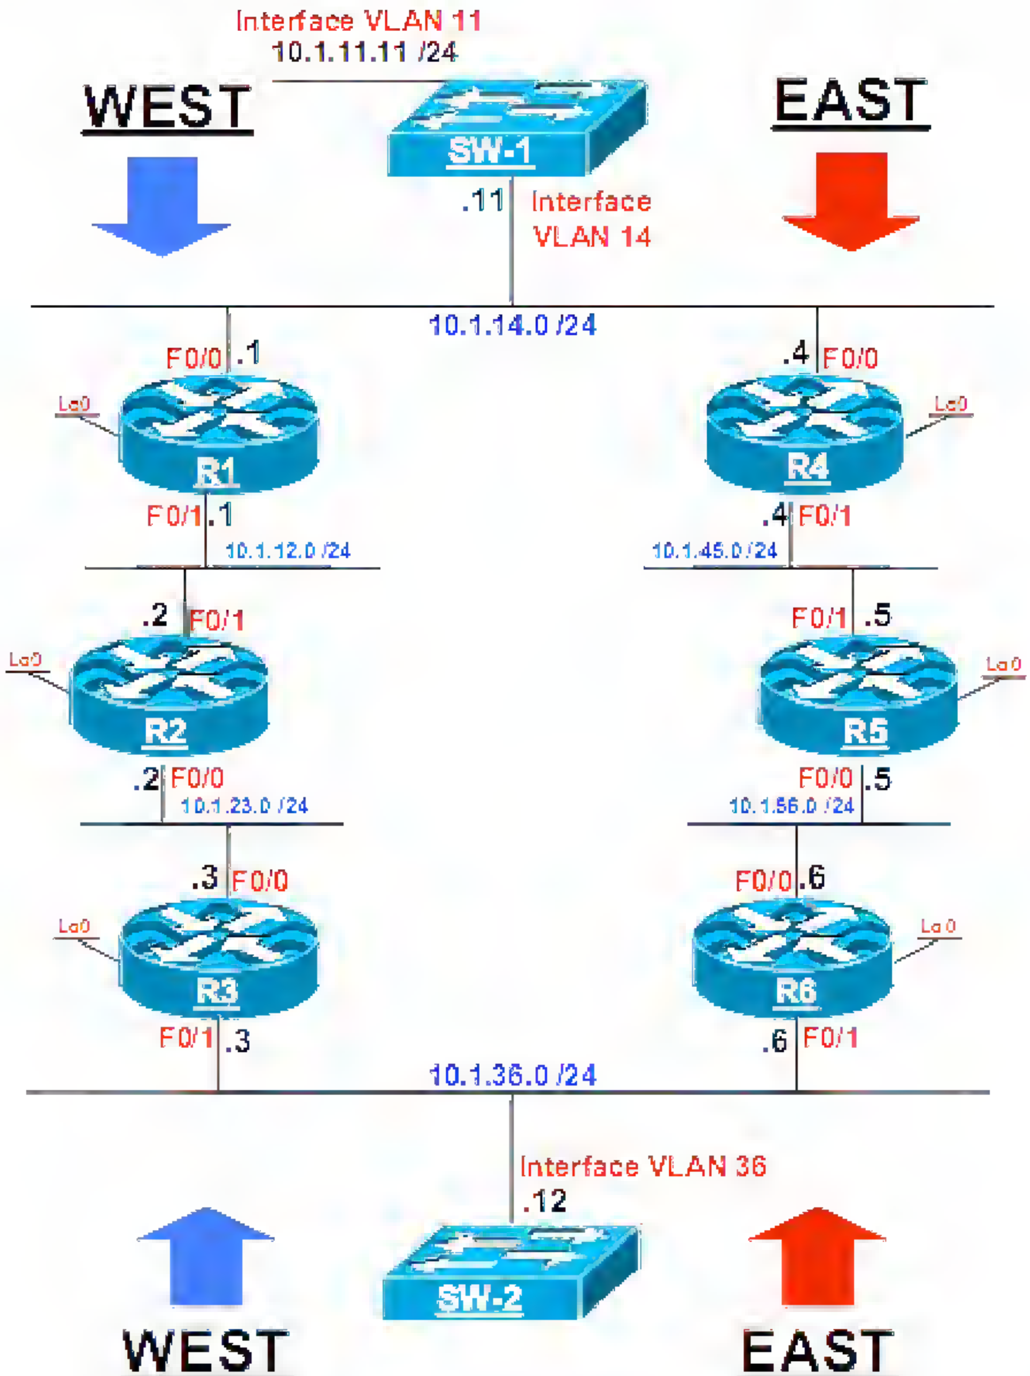

Lab Setup:

- Configure R1 and R4's F0/0 interface in VLAN 14
- Configure R1 and R2's F0/1 interface in VLAN 12
- Configure R2 and R3's F0/0 interface in VLAN 23
- Configure R3 and R6's F0/1 interface in VLAN 36
- Configure R5 and R6's F0/0 interface in VLAN 56
- Configure R4 and R5's F0/1 interface in VLAN 45
- Use the following IP addressing chart:

IP addressing Chart:

Router	VLAN	Interface & IP Addressing
R1	14 12	F0/0 – 10.1.14.1 /24 F0/1 – 10.1.12.1 /24 Lo0 – 10.1.99.1 /32
R2	12 23	F0/1 – 10.1.12.2 /24 F0/0 – 10.1.23.2 /24 Lo0 – 10.1.99.2 /32
R3	23 36	F0/0 – 10.1.23.3 /24 F0/1 – 10.1.36.3 /24 Lo0 – 10.1.99.3 /32
R4	14 45	F0/0 – 10.1.14.4 /24 F0/1 – 10.1.45.4 /24 Lo0 – 10.1.99.4 /32
R5	45 56	F0/1 – 10.1.45.5 /24 F0/0 – 10.1.56.5 /24 Lo0 – 10.1.99.5 /32
R6	56 36	F0/0 – 10.1.56.6 /24 F0/1 – 10.1.36.6 /24 Lo0 – 10.1.99.6 /32
SW1	11 14	Vlan11 – 10.1.11.11 /24 Vlan14 – 10.1.14.11 /24
SW2	36	Vlan36 – 10.1.36.12 /24

Task 1

Configure these devices such that when SW2, R2 or R5 communicate with VLAN 11 they should use the following policies:

- The traffic between SW2, R2 or R5 and Vlan 11 must always take the East path, even if the West path is closer.

- The **East path** should be the primary path, ensure that the solution covers any broken reachability on this path; the West path should be configured to be the backup path.
- One static route is allowed to be configured on R1, R4 and/or SW1

OSPF:

- VLAN 23 should be configured in area 23
- VLAN 12 should be configured in area 12
- VLAN 56 should be configured in area 56
- VLAN 45 should be configured in area 0
- Type 5 LSAs are NOT permitted in area 56
- Place loopback interfaces in the lower area IDs when possible.

Eigrp:

- ONLY configure VLAN 36 in Eigrp AS 100
- R6 should be configured as a stub router

DO NOT use GRE, IPsec and/or IPnIP tunnels to accomplish this task

Since this lab is one huge task, it requires lots of planning and design, this lab needs to be followed step by step because it is simply designed to teach some important concepts when redistribution is configured.

The first step is to enable routing on the switches; by default, IP routing is NOT enabled:

On SW1

SWx(config)#**ip routing**

In this lab all devices need to communicate with VLAN 11, therefore, they need to have reachability to VLAN 11, this reachability will be provided in the later solution/s, but VLAN 11 also needs to have a return path to these networks/subnets; since the East side must be the primary path and the West side needs to be the

backup, there must be a way to control this condition, this solution configured HSRP to accomplish this task.

The first step in accomplishing this task is to configure the static routes:

On R1 and R4

R1(config)#**ip route 10.1.11.0 255.255.255.0 10.1.14.11**

To verify the configuration:

On R1

R1#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

On R4

R4#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

Next, we must configure a static route on SW1, but what should be the next hop for this static route?
Should it be R4? Well even though R4 is in the East side (Primary path) but what if one of the links is down? How will the path switch over when one of the links in the primary path is down?
One way to resolve this issue is to use HSRP, R4 should be configured as the active router because it belongs to the East side, and R1 should be configured to be the standby router. Since no IP address is given for the Virtual IP, the VIP depicted for this solution is 10.1.14.100 (An arbitrary IP address).

On R1

R1(config)#int f0/0
R1(config-if)#standby 1 ip 10.1.14.100
R1(config-if)#standby 1 preempt

On R4

R4(config)#int f0/0
R4(config-if)#standby 1 ip 10.1.14.100
R4(config-if)#standby 1 preempt
R4(config-if)#standby 1 priority 110

To verify the configuration:

On R1

R1#show Standby brief

```
                     P indicates configured to preempt.
                     |
Interface   Grp  Prio  P  State    Active      Standby     Virtual IP
Fa0/0       1    100   P  Standby  10.1.14.4   local       10.1.14.100
```

On R4

R4#show Standby brief

```
                     P indicates configured to preempt.
                     |
Interface   Grp  Prio  P  State    Active      Standby     Virtual IP
Fa0/0       1    110   P  Active   Local       10.1.14.1   10.1.14.100
```

The next step is to configure the static route on SW1, this is configured for the return traffic for any of the networks/subnets in this topology:

On SW1

SW1(config)#**ip route** 0.0.0.0 0.0.0.0 10.1.14.100

To verify the configuration:

On SW1

SW1#Show ip route static

S* 0.0.0.0/0 [1/0] via 10.1.14.100

SW1#Ping 10.1.14.100

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.100, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
```

The next step is to configure IGP, this is done from top to bottom starting with OSPF area 0 in the East side, when configuring OSPF, the router ids are extremely important, unless the task specifically forbids it. Remember that the loopback interfaces must be configured in the lower area ID. Since area 56 is NOT allowed to have LSA type 5s, its configured as a stub.

On R4

```
R4(config)#router ospf 1
R4(config-router)#router-id 10.1.99.4
R4(config-router)#Netw 10.1.45.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.99.4 0.0.0.0 area 0
```

To verify the configuration:

On R4

R4#Show ip ospf interface brief

```
Interface   PID  Area        IP Address/Mask   Cost  State  Nbrs  F/C
Lo0         1    0           10.1.99.4/32      1     LOOP   0/0
Fa0/1       1    0           10.1.45.4/24      1     DR     0/0
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#router-id 10.1.99.5
R5(config-router)#Netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#Netw 10.1.56.5 0.0.0.0 area 56
R5(config-router)#Netw 10.1.99.5 0.0.0.0 area 0
R5(config-router)#area 56 stub
```

To verify the configuration:

On R5

R5#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O       10.1.99.4/32 [110/2] via 10.1.45.4, 00:03:06, FastEthernet0/1
```

The next step is to configure area 56:

On R6

```
R6(config)#router ospf 1
R6(config-router)#router-id 10.1.99.6
R6(config-router)#Netw 10.1.56.6 0.0.0.0 area 56
R6(config-router)#Netw 10.1.99.6 0.0.0.0 area 56
R6(config-router)#area 56 stub
```

To verify the configuration:

On R6

R6#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA    10.1.45.0/24 [110/2] via 10.1.56.5, 00:01:06, FastEthernet0/0
O IA    10.1.99.4/32 [110/3] via 10.1.56.5, 00:01:06, FastEthernet0/0
O IA    10.1.99.5/32 [110/2] via 10.1.56.5, 00:01:06, FastEthernet0/0
O*IA 0.0.0.0/0 [110/2] via 10.1.56.5, 00:00:17, FastEthernet0/0
```

To verify the configuration:

On R6

R6#Ping 10.1.99.4

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.99.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

R6#Ping 10.1.99.5

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.99.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note the reachability to VLAN 11 will be addressed later.

The next step is to configure the East side:

On R1

```
R1(config)#router ospf 1
R1(config-router)#router-id 10.1.99.1
R1(config-router)#Netw 10.1.12.1 0.0.0.0 area 12
R1(config-router)#Netw 10.1.99.1 0.0.0.0 area 12
```

To verify the configuration:

On R1

R1#Show ip ospf interface brief

```
Interface    PID   Area          IP Address/Mask      Cost  State   Nbrs  F/C
Lo0          1     12            10.1.99.1/32         1     LOOP    0/0
Fa0/1        1     12            10.1.12.1/24         1     DR      0/0
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#router-id 10.1.99.2
R2(config-router)#Netw 10.1.12.2 0.0.0.0 area 12
R2(config-router)#Netw 10.1.99.2 0.0.0.0 area 12
R2(config-router)#Netw 10.1.23.2 0.0.0.0 area 23
```

To verify the configuration:

On R2

R2#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O       10.1.99.1/32 [110/2] via 10.1.12.1, 00:01:28, FastEthernet0/1
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#router-id 10.1.99.3
R3(config-router)#Netw 10.1.23.3 0.0.0.0 area 23
```

R3(config-router)#Netw 10.1.99.3 0.0.0.0 area 23

To verify the configuration:

On R3

R3#Show ip route ospf

R3#

Note no routes are exchanged, next the routing table of R1 is examined:

On R1

R1#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.1.99.2/32 [110/2] via 10.1.12.2, 00:13:13, FastEthernet0/1
```

Note R1 does NOT have the R2's F0/0 interface in its routing table. We should examine R2's routing table:

On R2

R2#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O       10.1.99.1/32 [110/2] via 10.1.12.1, 00:15:51, FastEthernet0/1
O       10.1.99.3/32 [110/2] via 10.1.23.3, 00:11:57, FastEthernet0/0
```

Note R2 contains all routes in its routing table but does NOT advertise the routes from one area to another.
Remember ONLY an ABR can advertise routes from one area to another. But an ABR is a router that connects two different areas and this happens to be the case for R2, well….let's verify the fact that R2 is an ABR:

On R1

R1#Show ip ospf border

OSPF Process 1 internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

Note R1 does NOT see R2 as an ABR, this is because R2 does NOT have connectivity to area 0, to fix this problem a Loopback interface is created and placed in area 0 as follows:

<u>On R2</u>

R2(config)#int lo2
R2(config-if)#ip address 10.2.99.2 255.255.255.255
R2(config-if)#IP ospf 1 area 0

<u>To verify the configuration:</u>

<u>On R1</u>

<u>R1#Show ip ospf border</u>

OSPF Process 1 internal Routing Table

Codes: i - Intra-area route, I - Inter-area route
i 10.1.99.2 [1] via 10.1.12.2, FastEthernet0/1, ABR, Area 12, SPF 5

<u>R1#Show ip route ospf</u>

```
     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
O IA    10.1.23.0/24 [110/2] via 10.1.12.2, 00:02:10, FastEthernet0/1
O IA    10.2.99.2/32 [110/2] via 10.1.12.2, 00:02:10, FastEthernet0/1
O       10.1.99.2/32 [110/2] via 10.1.12.2, 00:02:10, FastEthernet0/1
O IA    10.1.99.3/32 [110/3] via 10.1.12.2, 00:02:10, FastEthernet0/1
```

<u>On R3</u>

<u>R3#Show ip route ospf</u>

```
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA    10.1.12.0/24 [110/2] via 10.1.23.2, 00:03:01, FastEthernet0/0
O IA    10.2.99.2/32 [110/2] via 10.1.23.2, 00:03:01, FastEthernet0/0
O IA    10.1.99.1/32 [110/3] via 10.1.23.2, 00:03:01, FastEthernet0/0
O IA    10.1.99.2/32 [110/2] via 10.1.23.2, 00:03:01, FastEthernet0/0
```

The next step is to configure Eigrp:

<u>On SW2</u>

SW2(config)#router eigrp 100
SW2(config-router)#no au
SW2(config-router)#Netw 10.1.36.12 0.0.0.0

To verify the configuration:

On SW2

SW2#Show ip eigrp interfaces

IP-EIGRP interfaces for process 100

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Vl36	0	0/0	0	0/10	0	0

On R3

R3(config)#router eigrp 100
R3(config-router)#no au
R3(config-router)#Netw 10.1.36.3 0.0.0.0

The next Eigrp router (R6) should also be configured as an Eigrp Stub router:

On R6

R6(config)#router eigrp 100
R6(config-router)#no au
R6(config-router)#Netw 10.1.36.6 0.0.0.0
R6(config-router)#Eigrp stub

To verify the configuration:

On R6

R6#Show ip eigrp neighbors

IP-EIGRP neighbors for process 100

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
1	10.1.36.12	Fa0/1	14	00:01:24	167	1002	0	3
0	10.1.36.3	Fa0/1	11	00:01:25	4	200	0	7

Since all the IGP configuration is completed, its time to start with redistribution. This configuration is performed starting from East side first and once again, the configuration is started from the top of the topology to the bottom:

On R4

On this router the static route is redistributed into OSPF, as follows:

```
R4(config)#router ospf 1
R4(config-router)#redistribute static subnets
```

To verify the configuration:

On R5

R5#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E2    10.1.11.0/24 [110/20] via 10.1.45.4, 00:00:52, FastEthernet0/1
O       10.1.99.4/32 [110/2] via 10.1.45.4, 00:00:52, FastEthernet0/1
O       10.1.99.6/32 [110/2] via 10.1.56.6, 01:01:46, FastEthernet0/0
```

R5#Ping 10.1.11.11

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R6

R6#Ping 10.1.11.11

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

The nest step is to redistribute Eigrp into OSPF, this must be done on R6:

On R6

```
R6(config)#router ospf 1
R6(config-router)#Redistribute Eigrp 100 Subnets
```

Note you should get the following error message:
Warning: Router is currently an ASBR while having only one area which is a stub area

The reason is that R6 is part of an OSPF stub area, and redistribution is NOT allowed in a stub area. To fix this problem, area 56 is converted to an NSSA area, as follows:

<u>On R5</u>

R5(config)#router ospf 1
R5(config-router)#No area 56 stub
R5(config-router)#**area 56 NSSA default-information-originate**

<u>On R6</u>

R6(config)#router ospf 1
R6(config-router)#no area 56 stub
R6(config-router)#**area 56 NSSA default-information-originate**

<u>To verify the configuration:</u>

<u>On R6</u>

<u>R6#Show ip route ospf</u>

```
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O IA    10.1.45.0/24 [110/2] via 10.1.56.5, 00:02:41, FastEthernet0/0
O IA    10.1.99.4/32 [110/3] via 10.1.56.5, 00:02:41, FastEthernet0/0
O IA    10.1.99.5/32 [110/2] via 10.1.56.5, 00:02:41, FastEthernet0/0
O*N2  0.0.0.0/0 [110/1] via 10.1.56.5, 00:02:41, FastEthernet0/0
```

<u>On R5</u>

<u>R5#Show ip route ospf</u>

```
   10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O E2   10.1.11.0/24 [110/20] via 10.1.45.4, 00:04:49, FastEthernet0/1
O N2   10.1.36.0/24 [110/20] via 10.1.56.6, 00:04:49, FastEthernet0/0
O      10.1.99.4/32 [110/2] via 10.1.45.4, 00:04:59, FastEthernet0/1
O      10.1.99.6/32 [110/2] via 10.1.56.6, 00:04:49, FastEthernet0/0
```

Note the "N2" route is there because Eigrp was redistributed into OSPF in the prior step, and the "E2" route is there because R5 is part of area 0 and R4 redistributed the static route into OSPF in one of the previous steps.

In the next step, OSPF is redistributed into Eigrp routing domain, as follows:

On R6

```
R6(config)#router eigrp 100
R6(config-router)#redistribute OSPF 1 metric 1 1 1 1 1
```

To verify the configuration:

On R6

R6#Show ip eigrp topology

```
IP-EIGRP Topology Table for AS(100)/ID(10.1.99.6)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 0.0.0.0/0, 1 successors, FD is 2560000256
       via Redistributed (2560000256/0)
P 10.1.45.0/24, 1 successors, FD is 2560000256
       via Redistributed (2560000256/0)
P 10.1.36.0/24, 1 successors, FD is 28160
       via Connected, FastEthernet0/1
P 10.1.56.0/24, 1 successors, FD is 2560000256
       via Redistributed (2560000256/0)
P 10.1.99.4/32, 1 successors, FD is 2560000256
       via Redistributed (2560000256/0)
P 10.1.99.5/32, 1 successors, FD is 2560000256
       via Redistributed (2560000256/0)
P 10.1.99.6/32, 1 successors, FD is 2560000256
       via Redistributed (2560000256/0)
```

Note the redistributed routes are marked as "Redistributed" in the Eigrp's topology table.

To Verify the configuration

On R3

R3#Show ip route Eigrp

```
R3#
```

Note R3 did NOT get the redistributed routes. This is because R6 is configured as "Eigrp stub", if no keyword is used with Eigrp stub configuration, the "connected" and the "summary" keyword will be added as default setting. The output of the following show command reveals that, as follows:

R6#Show run | S router eigrp 100

```
router eigrp 100
 redistribute ospf 1 metric 1 1 1 1 1
 network 10.1.36.6 0.0.0.0
 no auto-summary
 eigrp stub connected summary
```

To fix this problem the "Redistributed" keyword is configured, as follows:

On R6

R6(config)#Router eigrp 100
R6(config-router)#**Eigrp stub redistributed**

To verify the configuration:

On R3

R3#Show ip route Eigrp

```
     10.0.0.0/8 is variably subnetted, 12 subnets, 2 masks
D EX    10.1.45.0/24 [170/2560002816] via 10.1.36.6, 00:00:54, FastEthernet0/1
D EX    10.1.56.0/24 [170/2560002816] via 10.1.36.6, 00:00:54, FastEthernet0/1
D EX    10.1.99.4/32 [170/2560002816] via 10.1.36.6, 00:00:54, FastEthernet0/1
D EX    10.1.99.5/32 [170/2560002816] via 10.1.36.6, 00:00:54, FastEthernet0/1
D EX    10.1.99.6/32 [170/2560002816] via 10.1.36.6, 00:00:54, FastEthernet0/1
D*EX 0.0.0.0/0 [170/2560002816] via 10.1.36.6, 00:00:54, FastEthernet0/1
```

On R3

R3#Ping 10.1.11.11

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Note the East side is complete, its time to configure the West path, once again the approach will be from the top of the topology to the bottom; the first step is to redistribute the static route into OSPF, as follows:

<u>On R1</u>

```
R1(config)#router ospf 1
R1(config-router)#redistribute static subnets
```

<u>To verify the configuration:</u>

<u>On R2</u>

<u>R2#Show ip route ospf</u>

```
        10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O E2  10.1.11.0/24 [110/20] via 10.1.12.1, 00:01:03, FastEthernet0/1
O       10.1.99.1/32 [110/2] via 10.1.12.1, 00:01:03, FastEthernet0/1
O       10.1.99.3/32 [110/2] via 10.1.23.3, 06:46:40, FastEthernet0/0
```

Note the ping command to VLAN 11 will fail; this will be fixed in later steps.
In the next step a mutual redistribution is configured between OSPF and Eigrp on R3:

<u>First Eigrp is redistributed into OSPF, as follows:</u>

<u>On R3</u>

```
R3(config)#router ospf 1
R3(config-router)#redistribute eigrp 100 subnets
```

<u>To verify the configuration:</u>

<u>On R1</u>

<u>R1#Show ip route ospf</u>

```
     10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
O IA    10.1.23.0/24 [110/2] via 10.1.12.2, 00:06:54, FastEthernet0/1
O E2    10.1.45.0/24 [110/20] via 10.1.12.2, 00:01:15, FastEthernet0/1
O E2    10.1.36.0/24 [110/20] via 10.1.12.2, 00:01:15, FastEthernet0/1
O E2    10.1.56.0/24 [110/20] via 10.1.12.2, 00:01:15, FastEthernet0/1
O IA    10.2.99.2/32 [110/2] via 10.1.12.2, 00:06:54, FastEthernet0/1
```

O   10.1.99.2/32 [110/2] via 10.1.12.2, 00:06:54, FastEthernet0/1
O IA   10.1.99.3/32 [110/3] via 10.1.12.2, 00:06:54, FastEthernet0/1
O E2   10.1.99.4/32 [110/20] via 10.1.12.2, 00:01:15, FastEthernet0/1
O E2   10.1.99.5/32 [110/20] via 10.1.12.2, 00:01:15, FastEthernet0/1
O E2   10.1.99.6/32 [110/20] via 10.1.12.2, 00:01:15, FastEthernet0/1

Next, OSPF is redistributed into Eigrp:

On R3

R3(config)#Router eigrp 100
R3(config-router)#redistribute ospf 1 metric 1 1 1 1 1

To verify the configuration:

On SW2

SW2#Show ip route eigrp

10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
D EX   10.1.11.0/24 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.12.0/24 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.23.0/24 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.45.0/24 [170/2560000512] via 10.1.36.6, 00:18:42, Vlan36
D EX   10.1.56.0/24 [170/2560000512] via 10.1.36.6, 00:18:42, Vlan36
D EX   10.2.99.2/32 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.99.1/32 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.99.2/32 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.99.3/32 [170/2560000512] via 10.1.36.3, 00:01:49, Vlan36
D EX   10.1.99.4/32 [170/2560000512] via 10.1.36.6, 00:18:42, Vlan36
D EX   10.1.99.5/32 [170/2560000512] via 10.1.36.6, 00:18:42, Vlan36
D EX   10.1.99.6/32 [170/2560000512] via 10.1.36.6, 00:18:42, Vlan36
D*EX 0.0.0.0/0 [170/2560000512] via 10.1.36.6, 00:18:42, Vlan36

Note all the routes are there, the following verifies the path taken by SW2 to VLAN 11:

SW2#Show ip route 10.1.11.0

Routing entry for 10.1.11.0/24

Known via "eigrp 100", distance 170, metric 2560000512, type external
Redistributing via eigrp 100
Last update from 10.1.36.3 on Vlan36, 00:03:48 ago

```
Routing Descriptor Blocks:
* 10.1.36.3, from 10.1.36.3, 00:03:48 ago, via Vlan36
    Route metric is 2560000512, traffic share count is 1
    Total delay is 20 microseconds, minimum bandwidth is 1 Kbit
    Reliability 1/255, minimum MTU 1 bytes
    Loading 1/255, Hops 1
```

SW2 is receiving a specific route from the West side and a default route from the East side, therefore, SW2 prefers the path through the West side to VLAN 11. The following reveals this fact:

On R6

R6#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
O IA    10.1.45.0/24 [110/2] via 10.1.56.5, 00:55:33, FastEthernet0/0
O IA    10.1.99.4/32 [110/3] via 10.1.56.5, 00:55:33, FastEthernet0/0
O IA    10.1.99.5/32 [110/2] via 10.1.56.5, 00:55:33, FastEthernet0/0
O*N2 0.0.0.0/0 [110/1] via 10.1.56.5, 00:55:33, FastEthernet0/0
```

On R3

R3#Show ip route OSPF

```
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
O E2    10.1.11.0/24 [110/20] via 10.1.23.2, 00:22:04, FastEthernet0/0
O IA    10.1.12.0/24 [110/2] via 10.1.23.2, 00:22:04, FastEthernet0/0
O IA    10.2.99.2/32 [110/2] via 10.1.23.2, 00:22:04, FastEthernet0/0
O IA    10.1.99.1/32 [110/3] via 10.1.23.2, 00:22:04, FastEthernet0/0
O IA    10.1.99.2/32 [110/2] via 10.1.23.2, 00:22:04, FastEthernet0/0
```

Note SW2 takes the West path because ONLY a default route is given by R6 which was generated by R5 when the "NSSA" area was configured. There is no way to make R5 to send a more specific route to area 56, because R5 knows VLAN 11 as an LSA type 5, which is NEVER sent to a stub/NSSA area.

There may be few ways to accomplish this task, the following is what will be configured:

- **On R5 remove the F0/0 interface from OSPF process 1**
- **Configure the F0/0 interface of R5 in another OSPF process**
- **Redistribute between the new and the existing process**

Note after the above configuration R6 should have VLAN11 in its routing table.

On R5

```
R5(config)#router ospf 1
R5(config-router)#No Netw 10.1.56.5 0.0.0.0 area 56
R5(config-router)#NO area 56 nssa default
R5(config-router)#NO area 56 nssa
```

Note the adjacency will be dropped. The next step is to add the interface in another OSPF process, as follows:

```
R5(config)#router ospf 2
R5(config-router)#Netw 10.1.56.5 0.0.0.0 area 56
R5(config-router)#area 56 nssa
```

The last step is to redistribute between the two processes:

```
R5(config)#router ospf 1
R5(config-router)#redistribute OSPF 2 subnets

R5(config)#router ospf 2
R5(config-router)#redistribute OSPF 1 subnets
```

To verify the configuration:

On R6

R6#Show ip route ospf

```
     10.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
O N2    10.1.11.0/24 [110/20] via 10.1.56.5, 00:00:30, FastEthernet0/0
O N2    10.1.45.0/24 [110/1] via 10.1.56.5, 00:00:30, FastEthernet0/0
O N2    10.1.99.4/32 [110/2] via 10.1.56.5, 00:00:30, FastEthernet0/0
O N2    10.1.99.5/32 [110/1] via 10.1.56.5, 00:00:30, FastEthernet0/0
```

On R4

R4#Show ip route 10.1.36.0

```
Routing entry for 10.1.36.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 10.1.45.5 on FastEthernet0/1, 00:02:32 ago
  Routing Descriptor Blocks:
```

* 10.1.45.5, from 10.1.99.5, 00:02:32 ago, via FastEthernet0/1
Route metric is 20, traffic share count is 1

Note the output of the above show command reveals that R4 has reachability to VLAN 36. The last check should be the routing table of SW2:

SW2#Show ip route 10.1.11.0

Routing entry for 10.1.11.0/24
Known via "eigrp 100", distance 170, metric 2560000512, type external
Redistributing via eigrp 100
Last update from 10.1.36.6 on Vlan36, 00:04:28 ago
Routing Descriptor Blocks:
10.1.36.6, from 10.1.36.6, 00:04:28 ago, via Vlan36

Route metric is 2560000512, traffic share count is 1
Total delay is 20 microseconds, minimum bandwidth is 1 Kbit
Reliability 1/255, minimum MTU 1 bytes
Loading 1/255, Hops 1
*** 10.1.36.3, from 10.1.36.3, 00:04:28 ago, via Vlan36**
Route metric is 2560000512, traffic share count is 1
(The rest of the output is omitted)

Note SW2 has two routes to VLAN 11, one through R3 and the second one through R6, but the requirement of this task stated that SW2 should take the East path as primary, therefore, the metric of the path must be manipulated when performing the redistribution of OSPF into Eigrp, as follows:

On R6

Note R3 and R6 are both redistributing OSPF into Eigrp using the same metric parameters (1 1 1 1 1), in the following configuration, on R6, the bandwidth for OSPF redistributed routes are increased to 100, which makes the composite metric redistributed by R3 less attractive, hence, R6 should be the ONLY route in the routing table of SW2:

R6(config)#router eigrp 100
R6(config-router)#redistribute ospf 1 metric 100 1 1 1 1

To verify the configuration:

On SW2

SW2#Show ip route 10.1.11.0

Routing entry for 10.1.11.0/24
  Known via "eigrp 100", distance 170, metric 25600512, type external
  Redistributing via eigrp 100
  Last update from 10.1.36.6 on Vlan36, 00:08:40 ago
  Routing Descriptor Blocks:
  * 10.1.36.6, from 10.1.36.6, 00:08:40 ago, via Vlan36
      Route metric is 25600512, traffic share count is 1
      Total delay is 20 microseconds, minimum bandwidth is 100 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1

SW2#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms

SW2#Traceroute 10.1.11.11

Type escape sequence to abort.
Tracing the route to 10.1.11.11

  1 10.1.36.6 8 msec 0 msec 0 msec
  2 10.1.56.5 0 msec 0 msec 0 msec
  3 10.1.45.4 9 msec 0 msec 0 msec
  4 10.1.14.11 0 msec * 0 msec

Verifying R5's reachability to VLAN 11:

On R5

R5#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R5#Traceroute 10.1.11.11

Type escape sequence to abort.
Tracing the route to 10.1.11.11

  1 10.1.45.4 0 msec 0 msec 4 msec
  2 10.1.14.11 0 msec * 0 msec

To very R2's reachability to VLAN 11:

On R2

R2#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Traceroute 10.1.11.11

Type escape sequence to abort.
Tracing the route to 10.1.11.11

  1 10.1.12.1 4 msec 0 msec 4 msec
  2 10.1.14.11 4 msec * 0 msec

Note this router is taking the West path and needs to be changed.

On R2

R2#Show ip route 10.1.11.11

Routing entry for 10.1.11.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 10.1.12.1 on FastEthernet0/1, 03:14:17 ago
  Routing Descriptor Blocks:
  *** 10.1.12.1, from 10.1.99.1, 03:14:17 ago, via FastEthernet0/1**
      Route metric is 20, traffic share count is 1

On R3

R3#Show ip route 10.1.11.0

Routing entry for 10.1.11.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 2

  Redistributing via eigrp 100
  Advertised by eigrp 100 metric 1 1 1 1 1

```
Last update from 10.1.23.2 on FastEthernet0/0, 03:17:01 ago
Routing Descriptor Blocks:
* 10.1.23.2, from 10.1.99.1, 03:17:01 ago, via FastEthernet0/0
    Route metric is 20, traffic share count is 1
```

Note R3 is receiving an update from its Eigrp neighbor with an Administrative distance of 170 (Because its external), and its also receiving the same route from OSPF; in this case this router will prefer the path through OSPF, because through OSPF it has an administrative distance of 110.
In order to fix this problem the administrative distance of OSPF or Eigrp needs to be manipulated, the following solution manipulates the administrative distance of Eigrp external routes, as follows:

```
R3(config)#router eigrp 100
R3(config-router)#distance eigrp 90 109
```

Note if the administrative distance had to be manipulated in OSPF, the following command had to be used:

On R2

```
R3(config)#router ospf 1
R3(config-router)#distance ospf external 171
```

To verify the configuration:

On R3

R3#Show ip route 10.1.11.0

```
Routing entry for 10.1.11.0/24
  Known via "eigrp 100", distance 109, metric 25602816, type external
  Redistributing via ospf 1, eigrp 100
  Advertised by ospf 1 subnets
  Last update from 10.1.36.6 on FastEthernet0/1, 00:10:05 ago
  Routing Descriptor Blocks:
  * 10.1.36.6, from 10.1.36.6, 00:10:05 ago, via FastEthernet0/1
      Route metric is 25602816, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 100 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

Note the output of the above command reveals that R3 is taking the correct path, the following show command reveals the routing table of R2:

On R2

R2#Show ip route 10.1.11.11

```
Routing entry for 10.1.11.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 10.1.23.3 on FastEthernet0/0, 00:00:55 ago
  Routing Descriptor Blocks:
   10.1.23.3, from 10.1.99.3, 00:00:55 ago, via FastEthernet0/0
     Route metric is 20, traffic share count is 1
  * 10.1.12.1, from 10.1.99.1, 00:00:55 ago, via FastEthernet0/1
     Route metric is 20, traffic share count is 1
```

Note the output of the above command reveals that R2 has two paths to VLAN11, one through R1 and another through R3. The task states that this router should take the path through East side as primary and the path through the West side as backup.
One way to accomplish this task is to inject the route into R2's routing table as "E1". In OSPF, E1 routes are always preferred over E2.

On R3

```
R3(config)#router ospf 1
R3(config-router)#redistribute eigrp 100 subnets metric-type 1
```

To verify the configuration:

On R2

R2#Show ip route 10.1.11.0

```
Routing entry for 10.1.11.0/24
  Known via "ospf 1", distance 110, metric 21, type extern 1
  Last update from 10.1.23.3 on FastEthernet0/0, 00:01:43 ago
  Routing Descriptor Blocks:
  * 10.1.23.3, from 10.1.99.3, 00:01:43 ago, via FastEthernet0/0
     Route metric is 21, traffic share count is 1
```

R2#Ping 10.1.11.11

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

R2#Traceroute 10.1.11.11

Type escape sequence to abort.
Tracing the route to 10.1.11.11

```
 1 10.1.23.3 0 msec 4 msec 0 msec
 2 10.1.36.6 0 msec 4 msec 0 msec
 3 10.1.56.5 4 msec 4 msec 0 msec
 4 10.1.45.4 0 msec 4 msec 4 msec
 5 10.1.14.11 4 msec *  0 msec
```

Note R2 is now taking the East path to reach VLAN 11. The last item in the task stated **"ensure that the solution covers any broken reachability on this path".**

To test this condition, F0/1 interface of R5 is shutdown, as follows:

On R5

R5(config)#int f0/1
R5(config-if)#shutdown

To verify the configuration:

On R5

R5#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Traceroute 10.1.11.11

Type escape sequence to abort.
Tracing the route to 10.1.11.11

```
 1 10.1.56.6 0 msec 0 msec 0 msec
 2 10.1.36.3 0 msec 4 msec 0 msec

 3 10.1.23.2 0 msec 4 msec 4 msec
 4 10.1.12.1 0 msec 0 msec 4 msec
 5  *  *  *
 6  *  *  *
```

7 * *

Note the path is correct, but why is it NOT working?

To resolve this problem the routing table of SW1 is examined:

SW1#Show ip route static

S* 0.0.0.0/0 [1/0] via 10.1.14.100

The route is via the virtual IP address of HSRP, but remember that R4 is the active router and NOT R1, therefore, the HSRP should be configured to change the active router to be R1 when and if any of the links along the East path is broken.

To fix this problem:

- **An "object tracking" is configured**
- **R4 should be configured to monitor the route to VLAN 36, if the route is NOT available through the East side, it should lower its priority by 20 such that R1 becomes the active router.**

On R4

R4(config)#**track 1 ip route 10.1.36.0/24 reachability**

R4(config)#int f0/0
R4(config-if)#**standby 1 track 1 decrement 20**

Note you should get messages from HSRP stating that R1 is now the active router.

To verify the configuration:

On R4

R4#Show standby brief

```
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active      Standby     Virtual IP
Fa0/0       1    90  P Standby  10.1.14.1   local       10.1.14.100
```

On R5

R5#Ping 10.1.11.11

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R5#Traceroute 10.1.11.11

Type escape sequence to abort.
Tracing the route to 10.1.11.11

  1 10.1.56.6 4 msec 0 msec 0 msec
  2 10.1.36.3 4 msec 0 msec 4 msec
  3 10.1.23.2 4 msec 0 msec 0 msec
  4 10.1.12.1 4 msec 4 msec 4 msec
  5 10.1.14.11 4 msec *  0 msec
```

The last test that needs to be conducted is to enable the F0/1 interface of R5 and test the reachability, as follows:

On R5

```
R5(config)#int f0/1
R5(config-if)#no shut
```

To verify the configuration:

On R5

```
R5#Ping 10.1.11.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.11.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R5#Traceroute 10.1.11.11


Type escape sequence to abort.
Tracing the route to 10.1.11.11

  1 10.1.56.6 0 msec 4 msec 0 msec
  2 10.1.36.3 4 msec 4 msec 0 msec
  3 10.1.23.2 4 msec 4 msec 0 msec
```

4 10.1.12.1 0 msec 4 msec 4 msec
5 10.1.14.11 0 msec * 0 msec

Note this router is still taking the West path. Note for R4 to reach VLAN 11, it goes to 10.1.14.11, as follows:

R4#Sh ip route 10.1.11.0

```
Routing entry for 10.1.11.0/24
  Known via "static", distance 1, metric 0
  Redistributing via ospf 1
  Advertised by ospf 1 subnets
  Routing Descriptor Blocks:
  * 10.1.14.11
      Route metric is 0, traffic share count is 1
```

On R5, the route is redistributed via OSPF, but R5 receives a route to VLAN 11 from two different routing processes, as long as the path is equal, it will take the oldest route in its routing table. To fix this problem, R5 should be configured to always prefer OSPF process 1, as follows:

R5(config)#router ospf 1
R5(config-router)#**distance ospf external 109**

To verify the configuration:

On R5

R5#Sh ip route 10.1.11.0

```
Routing entry for 10.1.11.0/24
  Known via "ospf 1", distance 109, metric 20, type extern 2, forward metric 1
  Redistributing via ospf 2
  Advertised by ospf 2 subnets
  Last update from 10.1.45.4 on FastEthernet0/1, 00:01:15 ago
  Routing Descriptor Blocks:

  * 10.1.45.4, from 10.1.99.4, 00:01:15 ago, via FastEthernet0/1

      Route metric is 20, traffic share count is 1
```

R5#Traceroute 10.1.11.11

Type escape sequence to abort.

```
Tracing the route to 10.1.11.11

  1 10.1.45.4 0 msec 0 msec 4 msec
  2 10.1.14.11 4 msec *  0 msec
```

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 4 – Routing loops

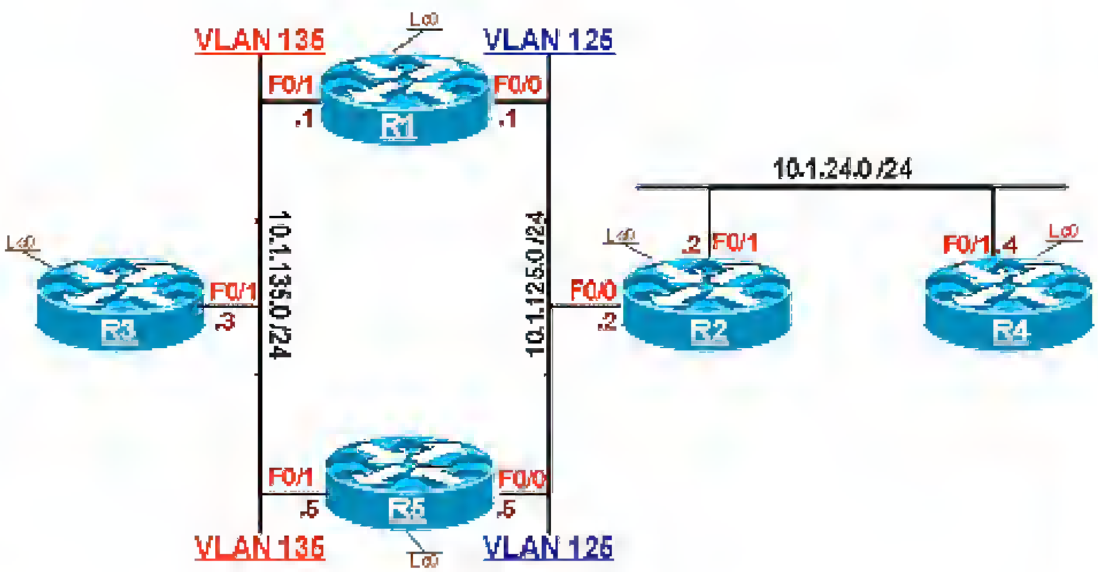

Lab Setup:

- Configure the rack according to the diagram.
- Use the IP addressing chart below for IP addressing scheme

IP addressing Chart:

Router	Interface / IP addressing
R1	F0/0 : 10.1.125.1/24 F0/1 : 10.1.135.1/24 Lo0: 10.1.99.1/32
R2	F0/0 : 10.1.125.2/24 F0/1 : 10.1.24.2/24 Lo0: 10.1.99.2/32
R3	F0/1 : 10.1.135.3/24 Lo0: 10.1.99.3/32
R4	F0/1 : 10.1.24.4/24 Lo0: 10.1.99.4/32
R5	F0/0 : 10.1.125.5/24 F0/1 : 10.1.135.5/24 Lo0: 10.1.99.5/32

Task 1

Configure RIPv2 on F0/1 interface of R2 and R4.

Advertise R4's Lo0 interface in this routing protocol. R4 should advertise this loopback interface such that R2 sees this loopback interface in its routing table with a prefix-length of /24 and a metric of 5.

On R2

```
R2(config)#router rip
R2(config-router)#no auto-summary
R2(config-router)#version 2
R2(config-router)#network 10.0.0.0
```

On R4

```
R4(config)#router rip
R4(config-router)#no auto-summary
R4(config-router)#version 2
R4(config-router)#network 10.0.0.0
```

To verify the configuration:

On R2

R2#Sh ip route rip

```
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
R       10.1.99.4/32 [120/1] via 10.1.24.4, 00:00:09, FastEthernet0/1
```

Note R4's loopback interface has a prefix-length of 32 with a metric of 1, but the task requested a prefix-length of 24 and a metric of 5.
This can be done in different ways, but in this solution an access-list, offset-list and route summarization is used to accomplish this task; an access-list is configured to identify the network, an offset-list will offset the hop count by 4 so when R2 receives the route it will have it as 5 hops and a route summarization is used to change the mask. As follows:

On R4

```
R4(config)#access-list 1 permit host 10.1.99.4

R4(config)#router rip
```

R4(config-router)#offset-list 1 out 4

Note when configuring an offset-list, the number after the "offset-list" keyword references an access-list number, if a "o" is configured instead of an access-list number, the offset-list will apply to all networks that R4 advertises. In this task this offset-list should only apply to the loopback 0 of R4, therefore, an access-list is used to be more specific.

To verify the configuration:

On R2

R2#Show ip route rip

```
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
R        10.1.99.4/32 [120/5] via 10.1.24.4, 00:00:12, FastEthernet0/1
```

Note the route has a metric of 5; the next step is to configure a summary route and summarize the route with a prefix-length of 24.

On R4

R4(config)#int F0/1
R4(config-if)#ip summary-address rip 10.1.99.0 255.255.255.0

To verify the configuration:

On R2

R2#Show ip route rip

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R        10.1.99.0/24 [120/1] via 10.1.24.4, 00:00:15, FastEthernet0/1
R        10.1.99.4/32 [120/5] via 10.1.24.4, 00:01:37, FastEthernet0/1
```

Note there are two routes, one with a prefix-length of 32 and the other with a prefix- length of 24. If the "Show ip route rip" command is entered in rapid succession, you will note that the route with the prefix-length of 32 is not being refreshed and the route will go into "Possibly Down" which means that the route is invalidated, as follows:

R2#Show ip route rip

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R        10.1.99.0/24 [120/1] via 10.1.24.4, 00:00:11, FastEthernet0/1
```

R   10.1.99.4/32 [120/5] via 10.1.24.4, 00:02:34, FastEthernet0/1

R2#Show ip route rip

     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R   10.1.99.0/24 [120/1] via 10.1.24.4, 00:00:14, FastEthernet0/1
R   10.1.99.4/32 is possibly down,
           routing via 10.1.24.4, FastEthernet0/1

On R2

R2#Show ip route rip

     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
R   10.1.99.0/24 [120/1] via 10.1.24.4, 00:00:15, FastEthernet0/1

Note the route with a prefix-length of 32 is no longer in the routing table and it is flushed out of the routing table.

To summarize:

- **By default, RIP sends updates every 30 seconds, that's why the timer for the route with a prefix-length of 24 was less then 30 seconds. It was getting refreshed every 30 seconds.**
- **By default, RIP's invalidation timer is 180 seconds, that's why the "possibly down" showed up in the routing table after about 3 minutes.**
- **By default, RIP's flush timer is 240 seconds, which is the total time it took to remove the route from the routing table.**

Task 2

Configure EIGRP 100 on VLAN 125. Advertise R1, R2 and R5's Lo0 interface in this routing protocol.

On R1

R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#Network 10.1.99.1 0.0.0.0
R1(config-router)#Network 10.1.125.1 0.0.0.0

On R2

R2(config)#router eigrp 100
R2(config-router)#no au

R2(config-router)#Network 10.1.99.2 0.0.0.0
R2(config-router)#Network 10.1.125.2 0.0.0.0

On R5

R5(config)#router eigrp 100
R5(config-router)#no au
R5(config-router)#Network 10.1.99.5 0.0.0.0
R5(config-router)#Network 10.1.125.5 0.0.0.0

To verify the configuration:

On R1

R1#Show ip route eigrp 100

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D       10.1.99.2/32 [90/409600] via 10.1.125.2, 00:03:44, FastEthernet0/0
D       10.1.99.5/32 [90/409600] via 10.1.125.5, 00:02:09, FastEthernet0/0
```

On R2

R2#Show ip route eigrp 100

```
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
D       10.1.99.1/32 [90/409600] via 10.1.125.1, 00:04:31, FastEthernet0/0
D       10.1.99.5/32 [90/409600] via 10.1.125.5, 00:02:55, FastEthernet0/0
```

On R5

R5#Show ip route eigrp 100

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D       10.1.99.1/32 [90/409600] via 10.1.125.1, 00:03:43, FastEthernet0/0
D       10.1.99.2/32 [90/409600] via 10.1.125.2, 00:03:41, FastEthernet0/0
```

Task 3

Configure OSPF area 0 on VLAN 135; also include R1, R3 and R5's Lo0 interface in this area.

On R1

R1(config)#router ospf 1
R1(config-router)#Netw 10.1.135.1 0.0.0.0 area 0
R1(config-router)#Netw 10.1.99.1 0.0.0.0 area 0

On R3

R3(config)#router ospf 1
R3(config-router)#Netw 10.1.135.3 0.0.0.0 area 0
R3(config-router)#Netw 10.1.99.3 0.0.0.0 area 0

On R5

R5(config)#router ospf 1
R5(config-router)#Netw 10.1.135.5 0.0.0.0 area 0
R5(config-router)#Netw 10.1.99.5 0.0.0.0 area 0

Note by now, all OSPF adjacencies should have come up, and a DR should be seen in the OSPF database with links to all routers in VLAN 135.

To verify the configuration:

On R1

R1#Show ip ospf database Network | b Attached

Attached Router: 10.1.99.1
Attached Router: 10.1.99.3
Attached Router: 10.1.99.5

Task 4

Configure mutual redistribution on R2 between EIGRP and RIP.

Before performing the redistribution, the routing table of R4 should be checked:

R4#Show ip route rip

```
     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
R       10.1.99.2/32 [120/1] via 10.1.24.2, 00:00:13, Ethernet0/1
R       10.1.125.0/24 [120/1] via 10.1.24.2, 00:00:13, Ethernet0/1
```

The next step, redistributing Eigrp 100 into RIP routing domain:

On R2

```
R2(config)#router rip
R2(config-router)#redistribute eigrp 100 metric 3
```

To verify the configuration:

On R4

Note the hop count is 1

R4#Show ip route rip

Note the hop is 3 as configured in the redistribution

```
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R       10.1.99.1/32 [120/3] via 10.1.24.2, 00:00:14, FastEthernet0/1
R       10.1.99.2/32 [120/1] via 10.1.24.2, 00:00:14, FastEthernet0/1
R       10.1.99.5/32 [120/3] via 10.1.24.2, 00:00:14, FastEthernet0/1
R       10.1.125.0/24 [120/1] via 10.1.24.2, 00:00:14, FastEthernet0/1
```

Note networks 10.1.125.0 /24 and 10.1.99.2 /32 have a hop count of 1, whereas, the other EIGRP routes have a hop count of 3.
The reason is that RIP on R2 can only be configured to include classful networks, meaning the entire 10.0.0.0. RIP includes all interfaces whose classless IP address falls inside of 10.0.0.0/8 network. Now what goes into the routing table is the best route, which is the route with the least number of hop counts.

In the Next step, RIP is redistributed into Eigrp:

On R2

```
R2(config)#router eigrp 100
R2(config-router)#redistribute rip metric 1 1 1 1 1
```

To verify the configuration:

On R5

R5#Show ip route eigrp 100

```
     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
D EX    10.1.24.0/24 [170/2560025856] via 10.1.125.2, 00:01:37, FastEthernet0/0
D EX    10.1.99.0/24 [170/2560025856] via 10.1.125.2, 00:01:37, FastEthernet0/0
D       10.1.99.1/32 [90/409600] via 10.1.125.1, 00:45:57, Ethernet0/0
D       10.1.99.2/32 [90/409600] via 10.1.125.2, 00:45:55, Ethernet0/0
```

Task 5

Configure mutual redistribution between Eigrp and OSPF on R1 and R5; ensure that the path from R3's Lo0 to R4's Lo0 is loop free and optimal.

The first step, performing a mutual redistribution between OSPF and Eigrp on R1 and R5:

On R1

R1(config)#router ospf 1
R1(config-router)#**redistribute eigrp 100 subnets**

R1(config)#router eigrp 100
R1(config-router)#**redistribute** ospf 1 **metric** 1 1 1 1 1

On R5

R5(config)#router ospf 1
R5(config-router)#**redistribute eigrp 100 subnets**

R5(config)#router eigrp 100
R5(config-router)#**redistribute** ospf 1 **metric** 1 1 1 1 1

Note:

From this point on, the routing tables on your lab/pod/rack might look different, as it is depended on the exact timing of your speed of entering the show/configurations. Routes pointing to R1 might be pointing to R5 on your lab/pod/rack and vice versa.
The routing table of R3 should be checked and the path towards R4's Lo0 (10.1.99.4/32) should be traced. Remember that R4's Lo0 is summarized to 10.1.99.0/24.

To verify the configuration:

On R3

R3#Show ip route 10.1.99.0

```
Routing entry for 10.1.99.0/24
 Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 10
 Last update from 10.1.135.1 on Ethernet0/1, 00:06:56 ago
 Routing Descriptor Blocks:
 * 10.1.135.1, from 10.1.99.1, 00:06:56 ago, via Ethernet0/1
     Route metric is 20, traffic share count is 1
```

From R3's routing table reveals that the route was learned from OSPF, the route type is E2 and the next-hop is R1.

You should "clear ip ospf proc" on R1, R3 and R5 before proceeding:

On R3

R3#Sh ip ospf neighbor

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.99.1         1   FULL/DR         00:00:39    10.1.135.1      FastEthernet0/1
10.1.99.5         1   FULL/BDR        00:00:39    10.1.135.5      FastEthernet0/1
```

R3#Ping 10.1.99.4 source 10.1.99.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.99.4, timeout is 2 seconds:
Packet sent with a source address of 10.1.99.3
.....
Success rate is 0 percent (0/5)
```

Note the above ping command is NOT successful even though the local router has a valid route for the destination.

The following examines the routing table of R1 for the destination, since R3's first hop is R1:

On R1

R1#Show ip route 10.1.99.0

Routing entry for 10.1.99.0/24
Known via "eigrp 100", distance 170, metric 2560025816, type external
Redistributing via eigrp 100, ospf 1
Advertised by ospf 1 subnets
Last update from 10.1.125.5 on FastEthernet0/0, 00:14:28 ago
Routing Descriptor Blocks:
10.1.125.5, from 10.1.125.5, 00:14:28 ago, via FastEthernet0/0

Route metric is 2560025816, traffic share count is 1
Total delay is 1010 microseconds, minimum bandwidth is 1 Kbit
Reliability 1/255, minimum MTU 1 bytes
Loading 1/255, Hops 1
* 10.1.125.2, from 10.1.125.2, 00:14:28 ago, via FastEthernet0/0
Route metric is 2560025816, traffic share count is 1
Total delay is 1010 microseconds, minimum bandwidth is 1 Kbit
Reliability 1/255, minimum MTU 1 bytes
Loading 1/255, Hops 1

Note from R1's perspective/view there are two ways to get to 10.1.99.0 /24, one via R2 and the other via R5 and both were learned from EIGRP.

But R1 should only have R2 as the next hop for this destination, what happened? Always remember, when mutual redistribution is performed on multiple points, at best we will have a suboptimal route and at worst we will end up with a routing loop. Since R3 can't ping R4's Lo0 even though the route is in it's routing table, this looks like a routing loop.....

Before moving on, you should ask yourself "why did all ICMP requests from R3 fail?" R1 has two paths to 10.1.99.0/24, one through R2 which is a valid path and the other through R5 which is NSV (Not So Valid). I guess this is another TLA (Three Letter Acronym) that can be added to the list (just a joke).

To verify this, the CEF decision making on R1 is examined:

On R1

(The output of the following show command is modified to ONLY display the specific entry)

R1#Show ip cef

Prefix	Next Hop	Interface
10.1.99.0/24	10.1.125.5	Ethernet0/0
	10.1.125.2	Ethernet0/0

Note the output of the above command reveals that there are two ways to get to 10.1.99.0

/24, let's be more specific, as follows:

R1#Show ip cef exact-route 10.1.99.3 10.1.99.4

10.1.99.3 -> 10.1.99.4 : Ethernet0/0 (next hop 10.1.125.5)

R1#Show ip cef exact-route 10.1.135.3 10.1.99.4

10.1.135.3 -> 10.1.99.4 : Ethernet0/0 (next hop 10.1.125.2)

Even though CEF has two equal cost routes, it is choosing only one route per src-dst pair.

Note the following reveals that if the Ping is sourced from R3 and NOT the loopback interface of R3, it is successful:

On R3

R3#Ping 10.1.99.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.99.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

To see the reason for the successful ping:

On R1

R1#Show ip cef exact-route 10.1.135.3 10.1.99.4

10.1.135.3 -> 10.1.99.4 : FastEthernet0/0 (next hop 10.1.125.2)

Note the output of the above command reveals that CEF has chosen R2 as its next hop towards 10.1.99.0/24.

What conclusion should be made from this example?
PING IS NOT ENOUGH TO TEST FOR ROUTING LOOPS

To verify the routing loops further:

On R3

R3#Traceroute 10.1.99.4 source 10.1.99.3

```
Type escape sequence to abort.
Tracing the route to 10.1.99.4

  1 10.1.135.1 20 msec 48 msec 20 msec
  2 10.1.125.5 20 msec 8 msec 28 msec
  3 10.1.135.1 24 msec 20 msec 20 msec
  4 10.1.125.5 20 msec 20 msec 48 msec
  5 10.1.135.1 16 msec 32 msec 36 msec
  6 10.1.125.5 32 msec 32 msec 88 msec
  7 10.1.135.1 36 msec 60 msec 28 msec
  8 10.1.125.5 84 msec 20 msec 44 msec
  9 10.1.135.1 92 msec 76 msec 88 msec
 10 10.1.125.5 80 msec 72 msec 80 msec
 11 10.1.135.1 68 msec 52 msec 36 msec
 12 10.1.125.5 144 msec 76 msec 76 msec
 13 10.1.135.1 108 msec 68 msec 92 msec
 14 10.1.125.5 56 msec 108 msec 96 msec
```

(The rest of the output is omitted)

I think this proves that we have a loop.

Its time to find the culprit of the routing loop.

Lets examine R5's routing table towards 10.1.99.0/24:

On R5

R5#Show ip route 10.1.99.0

```
Routing entry for 10.1.99.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 10
  Redistributing via eigrp 100
  Advertised by eigrp 100 metric 1 1 1 1 1
  Last update from 10.1.135.1 on Ethernet0/1, 01:06:33 ago
  Routing Descriptor Blocks:
  * 10.1.135.1, from 10.1.99.1, 01:06:33 ago, via Ethernet0/1
      Route metric is 20, traffic share count is 1
```

This doesn't look good at all. The route is learned from OSPF and the next hop is R1

Here is the flow of the events:

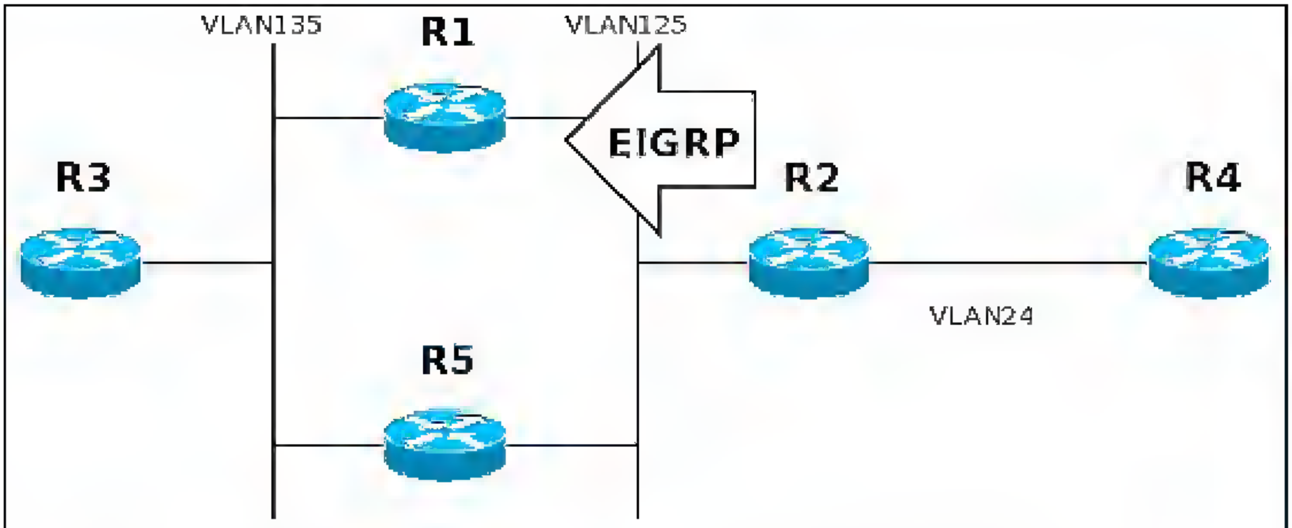

At first R2 sent an EIGRP update with the 10.1.99.0/24 network to R1 and R5.

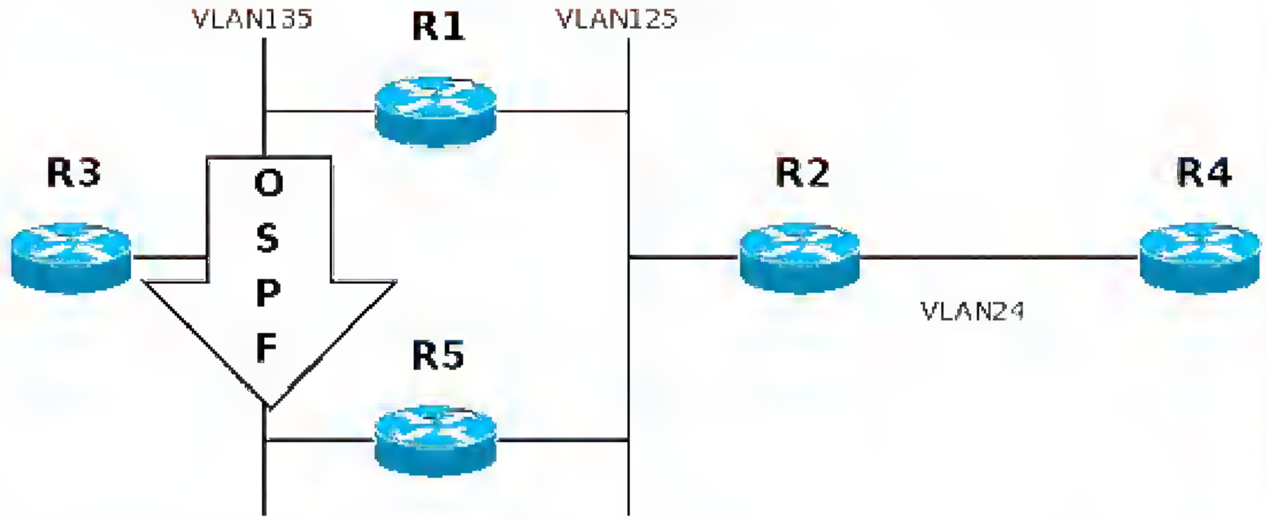

Then R1 redistributed the route into OSPF.

Now R5 has two routes to 10.1.99.0/24. One from OSPF and the other from EIGRP. R5 will choose the route with the lowest *administrative distance* (AD).

Here are R5's choices:

- Via R1 OSPF with AD 110
- Via R2 EIGRP with AD 170. 10.1.99.0/24 is an external EIGRP route, which was redistributed by R2 from RIP.

Since the route via OSPF has a lower AD, R5 will choose R1 as its next hop, via VLAN135.

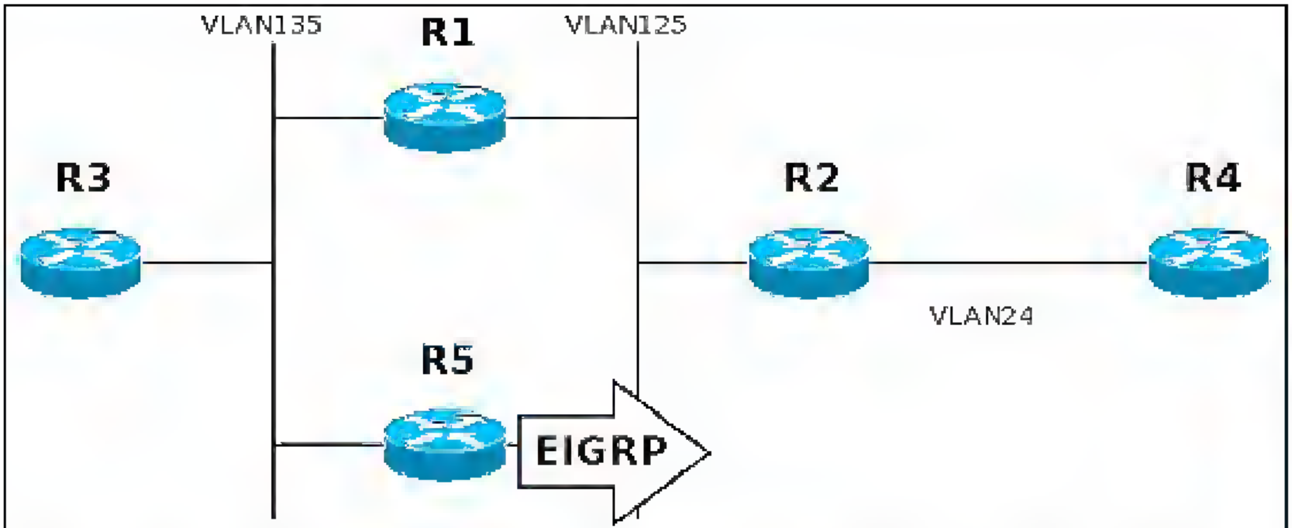

Now R5 redistributes the route from OSPF back to EIGRP. That's why R1 has a route to 10.1.99.0/24 from R5.

Does this process happen in the same way for R1? The answer is no. R1 will never get a 10.1.99.0/24 route from OSPF, as R5 will never redistribute it from EIGRP to OSPF. This is because redistribution happens only for routes which are actually installed in the routing table, and 10.1.99.0/24 route is installed as OSPF route on R5 and not as EIGRP. (remember the exception for connected routes)

How should we fix this? One way to fix this problem is to use the AD to our advantage. Since changing the AD of external EIGRP routes is not possible we will change the AD of OSPF routes. Here is the action plan for R5:

- Configure ACL to match all routes originated in EIGRP.
- Configure OSPF to raise the AD for routes matching the ACL

On R5

First the ACL:

```
R5(config)#access-list 1 permit 10.1.99.1
R5(config)#access-list 1 permit 10.1.99.2
R5(config)#access-list 1 permit 10.1.99.5
R5(config)#access-list 1 permit 10.1.99.0
R5(config)#access-list 1 permit 10.1.125.0
```

Next, let's change the distance:

```
R5(config)#router ospf 1
R5(config-router)#distance 171 0.0.0.0 255.255.255.255 1
```

The above command sets the distance for all the routes identified in the access-list 1 to 171. Note "0.0.0.0 255.255.255.255" was used to indicate that we don't care which router sends us the LSAs matching the ACL.

To verify the configuration:

On R5

R5#Sh ip route 10.1.99.0

```
Routing entry for 10.1.99.0/24
  Known via "eigrp 1", distance 170, metric 2560025816, type external
  Redistributing via eigrp 1, ospf 1
  Advertised by ospf 1 subnets
  Last update from 10.1.125.1 on FastEthernet0/0, 00:02:05 ago
  Routing Descriptor Blocks:
  * 10.1.125.2, from 10.1.125.2, 00:02:05 ago, via FastEthernet0/0
      Route metric is 2560025816, traffic share count is 1
      Total delay is 1010 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
    10.1.125.1, from 10.1.125.1, 00:02:05 ago, via FastEthernet0/0
      Route metric is 2560025816, traffic share count is 1
      Total delay is 1010 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

The output of the above show command reveals that there are two routes and both of them are via EIGRP, which is much better then the OSPF route.

To verify the configuration:

On R3

R3#Traceroute 10.1.99.4 source 10.1.99.3

```
Type escape sequence to abort.
Tracing the route to 10.1.99.4

  1 10.1.135.5 4 msec 0 msec 4 msec
  2 10.1.125.2 0 msec 0 msec 4 msec
  3 10.1.24.4 4 msec *  0 msec
```

From R3's Loopback 0 to R5, from R5 to R2 and from R2 to R4

R3#Ping 10.1.99.4 source 10.1.99.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.99.4, timeout is 2 seconds:
Packet sent with a source address of 10.1.99.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

The path from R1 to the Loopback 0 interface of R4 should also be checked, as follows:

On R1

R1#Sh ip route 10.1.99.0

```
Routing entry for 10.1.99.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Redistributing via eigrp 100
  Advertised by eigrp 100 metric 1 1 1 1 1
  Last update from 10.1.135.5 on FastEthernet0/1, 00:11:14 ago
  Routing Descriptor Blocks:
  * 10.1.135.5, from 10.1.99.5, 00:11:14 ago, via FastEthernet0/1
      Route metric is 20, traffic share count is 1
```

Now the opposite thing happens. R5 now has a route to 10.1.99.0 in its routing table, marked as Eigrp route and it is redistributing it into OSPF routing protocol.

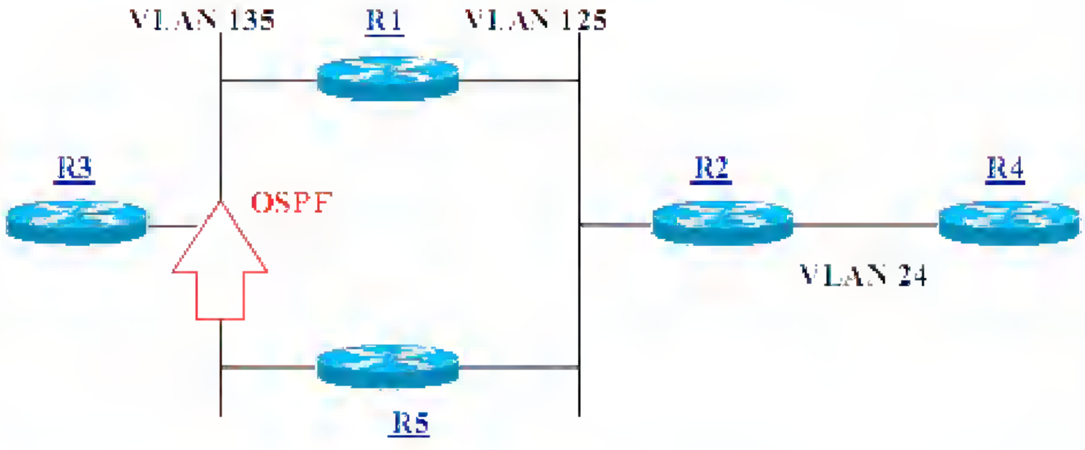

Note R1 receives the OSPF route and now R1 prefers the OSPF route since it has an AD of 110 versus EIGRP route from R2 with an AD of 170.
R1 then redistributes the route from OSPF into EIGRP.

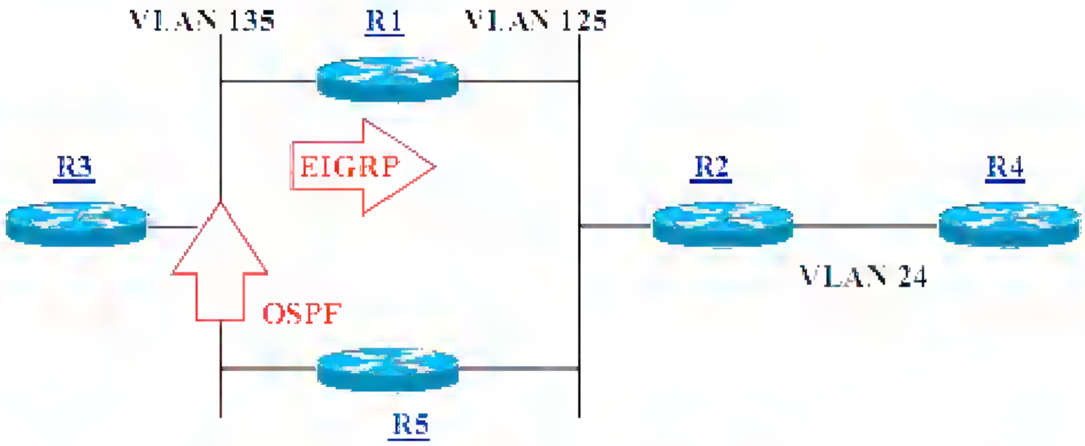

To fix this problem, AD is manipulated in the same manner as it was done on R5.

On R1:

An access-list is configured matching the EIGRP routes:

```
R1(config)#access-list 1 permit 10.1.99.2
R1(config)#access-list 1 permit 10.1.99.1
R1(config)#access-list 1 permit 10.1.99.0
R1(config)#access-list 1 permit 10.1.99.5
R1(config)#access-list 1 permit 10.1.125.0
```

The AD is changed for the routes identified in access-list 1:

```
R1(config-router)#distance 171 0.0.0.0 255.255.255.255 1
```

To verify the configuration:

On R1

R1#Show ip route 10.1.99.0

```
Routing entry for 10.1.99.0/24
  Known via "eigrp 1", distance 170, metric 2560025856, type external
  Redistributing via eigrp 1, ospf 1
  Advertised by ospf 1 subnets
  Last update from 10.1.125.2 on FastEthernet0/0, 00:01:20 ago
  Routing Descriptor Blocks:
  * 10.1.125.2, from 10.1.125.2, 00:01:20 ago, via FastEthernet0/0
      Route metric is 2560025856, traffic share count is 1
```

Total delay is 1010 microseconds, minimum bandwidth is 1 Kbit

Reliability 1/255, minimum MTU 1 bytes
Loading 1/255, Hops 1

The problem is fixed.

Task 6

Erase the startup config of all routers and reload them before proceeding to the next lab.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

IP
SLA

Lab 1 – IP SLA

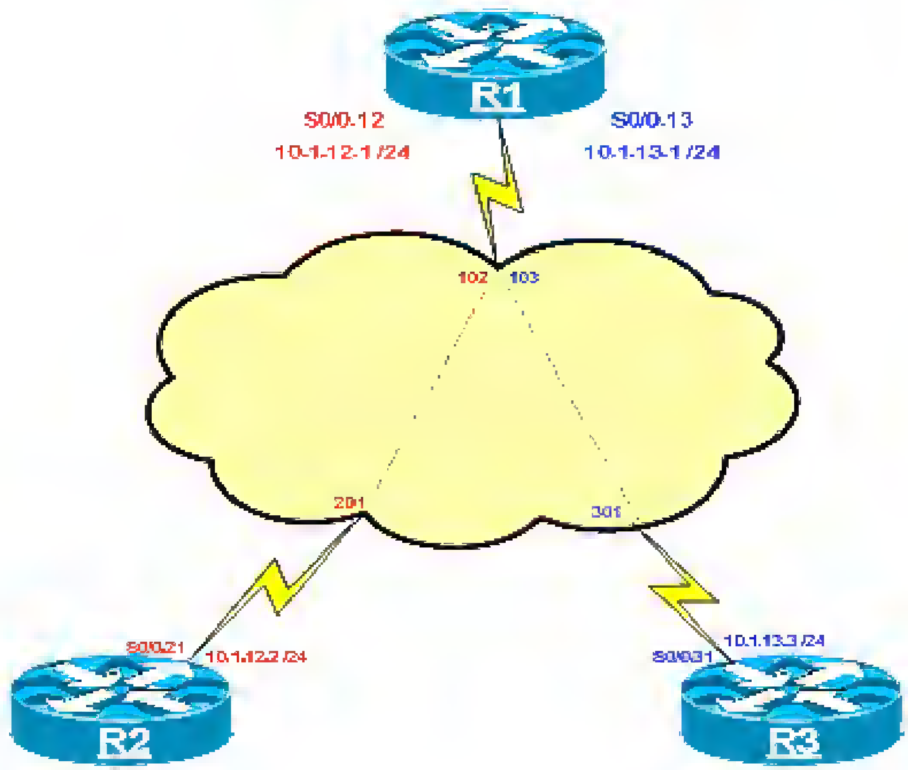

Lab Setup:

- Configure the the Frame-relay connection in a Point-to-point manner.
- Use the IP addressing chart below for IP address assignment.

IP addressing Chart:

Router	Interface / IP addressing
R1	S0/0.12 = 10.1.12.1 /24 S0/0.13 = 10.1.13.1 /24
R2	S0/0.21 = 10.1.12.2 /24

Task 1

Configure R1 and R2 so that they use UDP Echo packets to perform and determine the end-to-end response time. In this operation R1 should send these packets and R2 should be configured to respond to the packets with a time-stamp such that R1 can calculate the round trip time. This test should be performed for 30 seconds.

Using Cisco IP SLA, the performance of the network can be monitored; this can be performed without deploying a physical probe. A router can be configured to send a generated packet to the destination device and once the destination device receives this packet, the device will respond with time-stamp information for the source so the source can make the calculation on performance metric.
The UDP Echo operation measures end-to-end response time between a Cisco router and devices using IP.
In this task R2 should be configured as an IP SLA responder, a responder actually responds to Cisco IP SLA's request packets.

On R2

R2(config)#**ip sla monitor responder**

To verify the configuration:

On R2

R2#Sh ip sla monitor responder

IP SLAs Responder is: Enabled
Number of control message received: 0 Number of errors: 0
Recent sources:
Recent error sources:

Note the responder is enabled.

On R1

R1(config)#**IP SLA monitor 10**
R1(config-sla-monitor)#**type udpEcho dest-ipaddr 10.1.12.2 dest-port 12000**
R1(config-sla-monitor-udp)#**frequency 5**

Note the above commands configure a UDP ECHO to be sent to destination IP address of 10.1.12.2 (R2) to UDP port number 12000 every 5 seconds.
The following configures the scheduling parameters for the SLA operation to start immediately and continue for 30 seconds ONLY. Note the numeric value (10) after the "IP SLA monitor schedule"

command should match the number configured in the "IP SLA monitor" command above.

R1(config)#**IP SLA monitor schedule 10 start-time now life 30**

To test the configuration:

On R2

R2#Show ip sla monitor responder

IP SLAs Responder is: Enabled
Number of control message received: 5 Number of errors: 0
Recent sources:
    10.1.12.1 [05:33:19.611 UTC Tue Jan 15 2008]
    10.1.12.1 [05:33:14.611 UTC Tue Jan 15 2008]
    10.1.12.1 [05:33:09.611 UTC Tue Jan 15 2008]
    10.1.12.1 [05:33:04.611 UTC Tue Jan 15 2008]
    10.1.12.1 [05:32:59.611 UTC Tue Jan 15 2008]
Recent error sources:

On R1

The RTT time may vary in your test.

R1#Sh ip sla monitor statistics

Round trip time (RTT) Index 10
    Latest RTT: 30 ms
Latest operation start time: *05:27:24.176 UTC Mon Jan 14 2008
Latest operation return code: OK
Number of successes: 6
Number of failures: 0
peration time to live: 0

Task 2

Reconfigure the previous task to send packets with 1500 Bytes in size; this router should keep the statistics for a period of one hour

On R1

R1(config)#**NO ip sla monito 10**

R1(config)#**ip sla monitor** 10
R1(config-sla-monitor)#**type udpEcho dest-ipaddr 10.1.12.2 dest-port 12000**
R1(config-sla-monitor-udp)#frequency 5
R1(config-sla-monitor-udp)#**request-data-size** 1500
R1(config-sla-monitor-udp)#**hours-of-statistics-kept** 1

Note the "request-data-size" command can be used to set the packet size and the "hours-of-statistics-kept 1" command specifies that the stats should be kept for an hour ONLY.

Lastly the scheduling is invoked as follows:

R1(config)#**IP SLA monitor schedule** 10 **start-time now life 30**

To verify the configuration:

On R1

R1#Sh ip sla monitor statistics

Round trip time (RTT) Index 10
 Latest RTT: 773 ms
Latest operation start time: *06:04:34.600 UTC Mon Jan 14 2008
Latest operation return code: OK
Number of successes: 6
Number of failures: 0
Operation time to live: 0

Note the packet size was increased to 1500 bytes therefore, the RTT was affected, once again remember that your RTT may vary.

To verify the configuration:

On R1

R1#Show ip sla monitor configuration

SA Agent, Infrastructure Engine-II
Entry number: 10
Owner:
Tag:
Type of operation to perform: udpEcho
Target address: 10.1.12.2

```
Source address: 0.0.0.0
Target port: 12000
Source port: 0
Request size (ARR data portion): 1500
Operation timeout (milliseconds): 5000
Type Of Service parameters: 0x0
Verify data: No
Data pattern:
Vrf Name:
Control Packets: enabled
Operation frequency (seconds): 5
Next Scheduled Start Time: Start Time already passed
Group Scheduled : FALSE
Life (seconds): 30
Entry Ageout (seconds): never
Recurring (Starting Everyday): FALSE
Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 5000
Number of statistic hours kept: 1
Number of statistic distribution buckets kept: 1
Statistic distribution interval (milliseconds): 20
Enhanced History:
Number of history Lives kept: 0
Number of history Buckets kept: 15
History Filter Type: None
```

Task 3

Configure R3 to measure the response time taken to perform a TCP Connect operation between R3 and R1. R3 should be configured to generate TCP Connect messages, whereas, R1 should be configured such that it enhances the accuracy of the connection response time.

This task can be accomplished by configuring the IP SLAs TCP Connect operation, this operation is used to measure the response time taken to perform a TCP connect operation. To enhance the accuracy of the response time R1 should be configured as an IP SLA responder.

On R1

```
R1(config)#ip sla monitor responder
```

On R3

R3(config)#**ip sla monitor 30**
R3(config-sla-monitor)#**type TcpConnect dest-ipaddr 10.1.13.1 dest-port 23**
R3(config-sla-monitor-tcp)#**timeout 1000**
R3(config-sla-monitor-tcp)#**frequency 5**

R3(config)#**ip sla monitor schedule 30 life forever start-time now**

To verify the configuration:

On R1

R1#Sh ip sla monitor resp

IP SLA Monitor Responder is: Enabled
Number of control message received: 22 Number of errors: 0
Recent sources:
    10.1.13.3 [10:58:38.052 UTC Sat Jan 3 2009]
    10.1.13.3 [10:58:33.052 UTC Sat Jan 3 2009]
    10.1.13.3 [10:58:28.052 UTC Sat Jan 3 2009]
    10.1.13.3 [10:58:23.052 UTC Sat Jan 3 2009]
    10.1.13.3 [10:58:18.055 UTC Sat Jan 3 2009]
Recent error sources:

On R3

R3#Show ip sla monito statistics

Round trip time (RTT) Index 30
    Latest RTT: 36 ms
Latest operation start time: *16:27:28.141 UTC Tue Apr 16 2002
Latest operation return code: OK
Number of successes: 34
Number of failures: 0
Operation time to live: Forever

Task 4

Erase the startup config and reload the routers before proceeding to the next lab.

Lab – 2
Reliable Static Routing using IP SLA

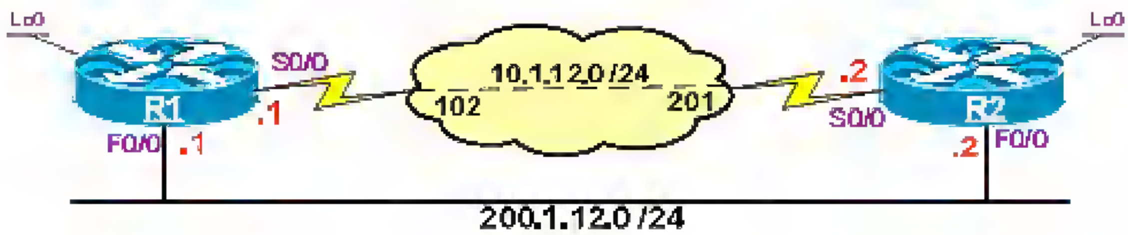

Lab Setup

- The frame-relay connection on R1 and R2 should be configured directly under the physical interface.
- The F0/0 interface of R1 and R2 should be configured in VLAN 12
- Run OSPF Area 0 on all interfaces of R1 and R2, Loopback interfaces should be advertised using their correct mask.
- Use the IP addressing chart below for IP addressing assignment

IP addressing Chart:

Router	Interface / IP addressing
R1	S0/0 = 10.1.12.1 /24 F0/0 = 200.1.12.1 /24 Lo0 = 1.1.1.1 /8
R2	S0/0 = 10.1.12.2 /24 F0/0 = 200.1.12.2 /24 Lo0 = 2.2.2.2 /8

Task 1

Configure two static routes on R1 to reach R2's loopback. The configuration should be such that if R1's frame-relay connection to R2 is reliably working, it should be the preferred path, but if R1 cannot reach R2 through the frame-relay cloud, R1 should take the path through its F0/0. DO NOT use EEK, backup interface, or PPP to accomplish this task.

To accomplish this task two floating static routes are configured as follows:

On R1

R1(config)#**ip route 2.0.0.0 255.0.0.0 10.1.12.2 50**
R1(config)#**ip route 2.0.0.0 255.0.0.0 200.1.12.2 100**

To verify the configuration:

On R1

R1#Show ip route | B Gate

```
Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
S    2.0.0.0/8 [50/0] via 10.1.12.2
C    200.1.12.0/24 is directly connected, FastEthernet0/0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0
```

To test the configuration:

On R2

To test this configuration, S0/0 interface of R2 is Shutdown and the routing table of R1 is checked, then, a Ping is generated from R1:

R2(config)#Int S0/0
R2(config-if)#**Shutdown**

On R1

R1#Show ip route | B Gate

Gateway of last resort is not set

```
C    1.0.0.0/8 is directly connected, Loopback0
S    2.0.0.0/8 [50/0] via 10.1.12.2
C    200.1.12.0/24 is directly connected, FastEthernet0/0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0
```

On R1

R1#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note this configuration did NOT accomplish the requirements of this task, because shutingdown the S0/0 interface of R2 did NOT effect R1 at ALL.

IP SLA ICMP ECHO can be used to monitor end-to-end response time between a Cisco router and another IP device, in this case, another Cisco router.
In the following configuration R1 is configured to generate "IpIcmpEcho" messages to the destination IP address of 10.1.12.2, the source IP address of these messages is set to 10.1.12.1. The timeout keyword specifies the amount of time an IP SLA operation waits for a response from its request packets. In this case the timeout is set to 500 milliseconds. The "frequency" keyword sets the rate at which the specified IP SLAs operation is repeated.

On R1

```
R1(config# ip sla monitor 1
R1(config-sla-monitor)#type echo protocol ipIcmpEcho 10.1.12.2 source-ipaddr 10.1.12.1
R1(config-sla-monitor-echo)#timeout 500
R1(config-sla-monitor-echo)#frequency 3
```

The above configuration is NOT enough for the router to generate the messages specified in the configuration, therefore, the router needs to be configured to start the above configuration operation.

The following configuration starts IP SLA operation 1, immediately with a life of the operation in this case set to forever.

```
R1(config)#ip sla monitor schedule 1 start-time now life forever
```

In the second last step of this configuration, the state of IP SLA operation is tracked for reachability:

R1(config)#track 2 rtr 1 reachability

The last step of this configuration, object tracking 2 is assigned to the primary static route:

R1(config)#NO ip route 2.0.0.0 255.0.0.0 10.1.12.2 50

R1(config)#ip route 2.0.0.0 255.0.0.0 10.1.12.2 50 track 2
R1(config)#ip route 2.0.0.0 255.0.0.0 200.1.12.2 100

Enable S0/0 interface of R2:

On R2

R2(config)#int S0/0
R2(config-if)#NO Shutdown

To verify the configuration:

On R1

R1#Show track 2

```
Track 2
  Response Time Reporter 1 reachability
  Reachability is Up
    1 changes, last change 00:00:04
  Latest operation return code: OK
  Latest RTT (millisecs) 39
  Tracked by:
    STATIC-IP-ROUTING 0
```

Note the rtr 1 is configured based on reachability, and the last operation was successful with a RTT of 39 ms.

R1#Show ip route | b Gateway

```
Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
S    2.0.0.0/8 [50/0] via 10.1.12.2
```

```
C    200.1.12.0/24 is directly connected, FastEthernet0/0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0
```

Sine the frame-relay link is up and network 2.0.0.0 /8 is reachable through the frame-relay cloud, its chosen as the best route.

R1#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/61/76 ms
```

To test the configuration:

On R2

To test the configuration, the S0/0 interface of R2 is shut down, as follows:

```
R2(config)#int S0/0
R2(config-if)#Shutdown
```

Note the shutting down the S0/0 (The Frame-relay) interface of R2 did not effect the Serial0/0 interface of R1:

On R1

R1#Sh ip int br | Inc Serial0/0_

```
Serial0/0              10.1.12.1       YES manual up                    up
```

To test the operation of IP SLA/tracking:

On R1

R1#Sh track 2

```
Track 2
  Response Time Reporter 1 reachability
  Reachability is Down
    3 changes, last change 00:09:19
  Latest operation return code: Timeout
```

Tracked by:
  STATIC-IP-ROUTING 0

Note because reachability is down (The IP SLA operation can NOT send IpIcmpEcho to 10.1.12.2), the static route to network 2.0.0.0 /8 with an administrative distance of 50 is removed and the static route with an Administrative distance of 100 is injected into the routing table.

R1#Show ip route | b Gateway

```
Gateway of last resort is not set
                                        Note the next hop IP address changed
C    1.0.0.0/8 is directly connected, Loopback0
S    2.0.0.0/8 [100/0] via 200.1.12.2
C    200.1.12.0/24 is directly connected, FastEthernet0/0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.12.0 is directly connected, Serial0/0
```

Note the following Traceroute & Ping command reveals that the reachability is now through the Fast Ethernet interface.

R1#traceroute 2.2.2.2

```
Type escape sequence to abort.
Tracing the route to 2.2.2.2

  1 200.1.12.2 20 msec *  48 msec
```

R1#Ping 2.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Note because of the floating static route and the IP SLA configuration, if the S0/0 interface of R2 is brought back up, the track 2 reachability will be UP, therefore, the primary static route will be injected back into the routing table and the backup static route will be removed.

To test the configuration:

On R2

```
R2(config)#int S0/0
R2(config-if)#NO shut
```

On R1

```
R1#Sh ip route | Inc 2.0.0.0

S    2.0.0.0/8 [50/0] via 10.1.12.2

R1#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 2

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab – 3
Reliable Conditional Default Route Injection Using IP SLA

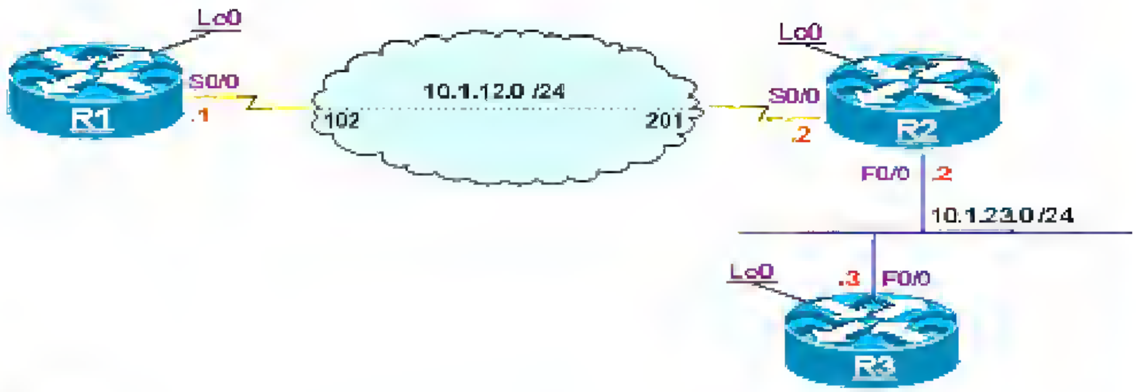

Lab Setup:

- The frame-relay connection on R1 and R2 should be configured directly under the physical interface.
- The F0/0 interface of R2 and R3 should be configured in VLAN 23
- Use the IP addressing chart below for IP addressing assignment

IP addressing Chart:

Router	Interface / IP addressing
R1	S0/0 = 10.1.12.1 /24
R2	S0/0 = 10.1.12.2 /24 F0/0 = 10.1.23.2 /24
R3	F0/0 = 10.1.23.3 /24 Lo0 = 3.3.3.3 /24

Task 1

Configure RIPv2 on the link that between R1 to R2 and OSPF area 0 between R2 and R3. R3 should run OSPF on all of it's directly connected interfaces.

On R1 & R2

```
(config-if)#router rip
(config-router)#no auto
(config-router)#ver 2
(config-router)#net 10.0.0.0
```

On R2

```
R2(config)#router rip
R2(config-router)#passive-interface F0/0

R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
```

To verify the configuration:

On R2

R2#Show ip route ospf | Inc O

```
O       3.3.3.3 [110/11] via 10.1.23.3, 00:07:02, FastEthernet0/0
```

On R1

R1#Show ip route rip | Inc R

```
R       10.1.23.0 [120/1] via 10.1.12.2, 00:00:26, Serial0/0
```

Task 2

Configure R2 to advertise a default route into OSPF routing domain. The default route should ONLY be injected if R2 and R1 have reachability through the frame-relay cloud. You are NOT allowed to use static route or IP SLA to accomplish this task.

Since the use of IP SLA and static route is prohibited, PPP is used to accomplish this task; when configuring PPP on any link, a host route is injected, therefore, the host route can be identified by an access-list, and the access-list is referenced in the route-map, and finaly the route-map is referenced by "default-information originate" router configuration command, as follows:

Step 1:
The following configures PPP on Frame-relay:

On R1 and R2

R2(config)#int S0/0
R2(config-if)#NO ip addr

On R1

R1(config)#**Int virtual-template** 12
R1(config-if)#ip addr 10.1.12.1 255.255.255.0

R1(config-if)#Int S0/0
R1(config-if)#**frame-relay interface-dlci** 102 **ppp virtual-template** 12

On R2

R2(config)#**Int virtual-template** 21
R2(config-if)#ip addr 10.1.12.2 255.255.255.0

R2(config-if)#Int S0/0
R2(config-if)#**frame-relay interface-dlci** 201 **ppp virtual-template** 21

To verify the configuration:

On R1

Note the output of the following Show command reveals the host route that is injected by PPP:

R1#Sh ip route | Inc C

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
C 10.1.12.2/32 is directly connected, Virtual-Access2
C 10.1.12.0/24 is directly connected, Virtual-Access2

On R2

R2#Sh ip route | Inc C

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
C 10.1.12.1/32 is directly connected, Virtual-Access2
C 10.1.12.0/24 is directly connected, Virtual-Access2
C 10.1.23.0/24 is directly connected, FastEthernet0/0

Note the two routers are exchanging routes:

On R1

R1#Sho ip route rip

10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R 10.1.23.0/24 [120/1] via 10.1.12.2, 00:00:25, Virtual-Access2

Step 2:
An access-list is configured to reference the host route generated by PPP:

R2(config)#**access-list 1 permit host 10.1.12.1**

Step 3:
A route-map is configured to reference the access-list:

R2(config)#route-map **TST** permit 10
R2(config-route-map)#match ip addr 1

Step 4:
In this final step a "default-information originate" is configured referencing the route-map:

R2(config-route-map)#router ospf 1
R2(config-router)#**default-information originate route-map TST**

To verify the configuration:

On R3

R3#Show ip route ospf | Inc O

O*E2 0.0.0.0/0 [110/1] via 10.1.23.2, 00:00:51, FastEthernet0/0

To test the configuration:

Note once R1's S0/0 goes down, the host route is removed and the condition of the route-map TST is no longer true, therefore, the default route is removed.

R1(config)#int S0/0
R1(config-if)#**Shut**

On R3

R3#Show ip route ospf | Inc O
R3#

To test this condition further:

On R1

R1(config)#int S0/0
R1(config-if)#No shut

On R3

R3#Show ip route ospf | Inc O

O*E2 0.0.0.0/0 [110/1] via 10.1.23.2, 00:00:08, FastEthernet0/0

Task 3

Re-configure R2 to advertise a default route into OSPF routing domain. The default route should ONLY be injected if R2 and R1 have reachability through the frame-relay cloud. You should use IP SLA to accomplish this task.

On R2

Step one:
The access-list, route-map and the default-information originate commands are removed:

R2(config)#**NO access-list 1**
R2(config)#**NO route-map TST**

R2(config)#router ospf 1
R2(config-router)#**NO default-information originate**

Configuring IP SLA to generate IP ICMP Echo messages, the timeout and frequency can be set to any value:

R2(config)#**ip sla monitor 10**
R2(config-sla-monitor)#**type echo protocol ipicmpEcho 10.1.12.1 source-ipaddr 10.1.12.2**
R2(config-sla-monitor-echo)#**timeout 250**
R2(config-sla-monitor-echo)#**frequency 5**

Note even though the IP SLA is configured, it won't start unless its configured to do so; when starting the operation, the start-time and life of these messages are defined:

R2(config)#**ip sla monitor schedule 10 start-time now life forever**

The IP SLA operation is tracked in track 2 for reachability:

R2(config)#**track 2 rtr 10 reachability**

Since the track is referenced in the following default route, if the IP SLA operation fails, track 2 will go down and the default route is removed. Remember, in order to initiate a default route in OSPF, the local router must have a default route or else it won't generate a default route:

R2(config)#**ip route 0.0.0.0 0.0.0.0 null0 track 2**

R2(config)#router ospf 1
R2(config-router)#**default-information originate**

To verify the configuration:

On R3

R3#Show ip route ospf

O*E2 0.0.0.0/0 [110/1] via 10.1.23.2, 00:17:38, FastEthernet0/0

To test the configuration:

On R1

R2#Show track 2

```
Track 2
  Response Time Reporter 10 reachability
  Reachability is Up
    2 changes, last change 00:00:09
  Latest operation return code: OK
  Latest RTT (millisecs) 36
  Tracked by:
    STATIC-IP-ROUTING 0
```

Note reachability is UP, therefore, the default route should be present in R3's routing table:

On R3

R3#Show ip route ospf

```
O*E2 0.0.0.0/0 [110/1] via 10.1.23.2, 00:02:06, FastEthernet0/0
```

To test this configuration:

1. **Serial0/0 interface of R1 is Shutdown.**
2. **The state of track 2 is checked on R2, the state must be Down.**
3. **The routing table of R3 is checked, it should NOT have a default route.**
4. **Serial0/0 interface of R1 is enabled.**
5. **The state of track 2 is checked on R2, the state must be UP.**
6. **The routing table of R3 is checked, it should have a default route.**

Step one:

On R1

```
R1(config)#Int S0/0
R1(config-if)#Shut
```

Step Two:

On R2

R2#Show track 2

Track 2
  Response Time Reporter 10 reachability
  Reachability is Down
    4 changes, last change 00:04:25
  Latest operation return code: Timeout
  Tracked by:
    STATIC-IP-ROUTING 0

Step Three:

On R3

R3#Show ip route ospf

R3#

Step Four:

On R1

R1(config)#Int S0/0
R1(config-if)#No **shut**

Step Five:

R2#Show track 2

Track 2
  Response Time Reporter 10 reachability
  Reachability is Up
    5 changes, last change 00:00:22
  Latest operation return code: OK
  Latest RTT (millisecs) 35
  Tracked by:
    STATIC-IP-ROUTING 0

Step Six:

R3#Show ip route ospf

O*E2 0.0.0.0/0 [110/1] via 10.1.23.2, 00:01:11, FastEthernet0/0

Task 4

Configure R2 to advertise a default route into RIPv2 routing domain. The default route should ONLY be injected if R2 has reachability to R3 through the switched connection. You should use IP SLA to accomplish this task.

Note R2 will NOT be aware if the F0/0 interface of R3 goes down, therefore, if the F0/0 interface of R3 is down, the F0/0 interface of R2 will remain in UP/UP state.

To configure this injection of default route reliably, once again, the IP SLA operation is configured like the previous tasks, but the difference in this configuration is the following:

A fake static route is created so it can be utilized to accomplish this task, this static route can be for any network, this network does NOT exist; this static route is tracked by the IP SLA operation and referenced in a route-map, the route-map is referenced in the "default-information originate" router configuration command, therefore, if R2 fails to reach R3's F0/0 IP address through the IP SLA operation, this static route is removed, if the static route is removed the condition of the route-map will NOT be true, therefore, the default route is removed.

On R2

The following creates a fake static route, in this case 3.3.3.3 /32 IP address is chosen:

R2(config)#**ip route 3.3.3.3 255.255.255.255 null0**

The following access-list is created to identify the fake static route:

R2(config)#**access-list 1 permit host 3.3.3.3**

A route-map is configured and access-list 1 is referenced:

R2(config)#route-map TST permit 10
R2(config-route-map)#**match ip address 1**

The following configuration instructs the router to inject a default route ONLY if the condition of the route-map is true; the condition of the route-map can only be true if 3.3.3.3 /32 exists:

R2(config)#router rip
R2(config-router)#**default-information originate** route-map TST

To verify the configuration:

On R1

Note the default route is injected:

R1#Show ip route rip | Inc R

R 10.1.23.0 [120/1] via 10.1.12.2, 00:00:01, Serial0/0
R* 0.0.0.0/0 [120/1] via 10.1.12.2, 00:00:01, Serial0/0

To test the configuration:

On R3

The F0/0 interface of R3 is Shutdown:

R3(config)#int f0/0
R3(config-if)#**Shut**

Note even though the F0/0 interface of R3 is in shutdown mode, the default route is still injected, as follows:

R1#Sh ip route rip | Inc R

R 10.1.23.0 [120/1] via 10.1.12.2, 00:00:17, Serial0/0
R* 0.0.0.0/0 [120/1] via 10.1.12.2, 00:00:17, Serial0/0

To inject a reliable default route, an IP SLA Monitor is configured to track the reachability of R3's F0/0 interface, this is called a reliable conditional default gateway injection, as follows:

R2(config)#ip sla monitor 10
R2(config-sla-monitor)#**type echo protocol IpIcmpEcho 10.1.23.3 source-ipaddr 10.1.23.2**
R2(config-sla-monitor-echo)#**timeout 250**
R2(config-sla-monitor-echo)#**frequency 3**

R2(config)#**ip sla monitor schedule 1 start-time now life forever**

R2(config)#**track 1 rtr 10 reachability**

The following command tracks the static route created earlier:

R2(config)#**NO ip route 3.3.3.3 255.255.255.255 null 0**

R2(config)#**ip route 3.3.3.3 255.255.255.255 null 0 track 1**

The F0/0 interface of R3 is re-enabled:

```
R3(config)#int f0/0
R3(config-if)#NO Shut
```

To verify the configuration:

On R2

R2#Sh track

```
Track 10
  Response Time Reporter 1 reachability
  Reachability is Up
    2 changes, last change 00:00:06
  Latest operation return code: OK
  Latest RTT (millisecs) 1
  Tracked by:
    STATIC-IP-ROUTING 0
```

R1#Show ip route rip | Inc R

```
R       10.1.23.0 [120/1] via 10.1.12.2, 00:00:16, Serial0/0
R*      0.0.0.0/0 [120/1] via 10.1.12.2, 00:00:16, Serial0/0
```

To test the configuration:

The following is how the test will be conducted:

- F0/0 interface of R3 is shutdown.
- A "Debug Track" on R2 and "Debug ip icmp" on R3 is configured.
- The routing table of R1 is checked; if the configuration was performed properly, R2 should remove the fake static route to 3.3.3.3/32, once this happens, the condition of the route-map (TST) is no longer true, therefore, the default route is removed and it will NOT be in the routing table of R1.
- The F0/0 interface of R3 is enabled (No Shut), if the configuration was performed properly, R2 should inject the default route back into RIP routing domain.

On R2

R2#Debug track

On R3

R3#Debug ip icmp

The interface is shutdown:

R3(config)#int F0/0
R3(config-if)#**shut**

Note on R2 you should receive the following messages:

00:46:04.507: Track: 1 Change #1 rtr 1, **reachability Up->Down**
00:46:04.511: %TRACKING-5-STATE: 1 **rtr 1 reachability Up->Down**

The routing table of R1 is verified:

On R1

R1#Show ip route rip | Inc R

R   10.1.23.0 [120/1] via 10.1.12.2, 00:00:09, Serial0/0

The output of the above command reveals that R1 no longer has the default route in its routing table.

The F0/0 interface of R3 is brought back up:

On R3

R3(config)#int f0/0
R3(config-if)#No **Shut**

On R2

You should receive the following message:

11:37:24.972: Track: 1 **Change #2 rtr 10, reachability Down->Up**

R2#Sh track

Track 10
  Response Time Reporter 1 reachability
  Reachability is Up
    2 changes, last change 00:02:29
  Latest operation return code: OK
  Latest RTT (millisecs) 1

Tracked by:
  STATIC-IP-ROUTING 0

The routing table of R1 is checked:

R1#sh ip route rip | Inc R

```
R      10.1.23.0 [120/1] via 10.1.12.2, 00:00:16, Serial0/0
R*     0.0.0.0/0 [120/1] via 10.1.12.2, 00:00:16, Serial0/0
```

Note the default route is in the routing table of R1. This may take few seconds.

Task 5

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab – 4
Object Tracking in HSRP Using IP SLA

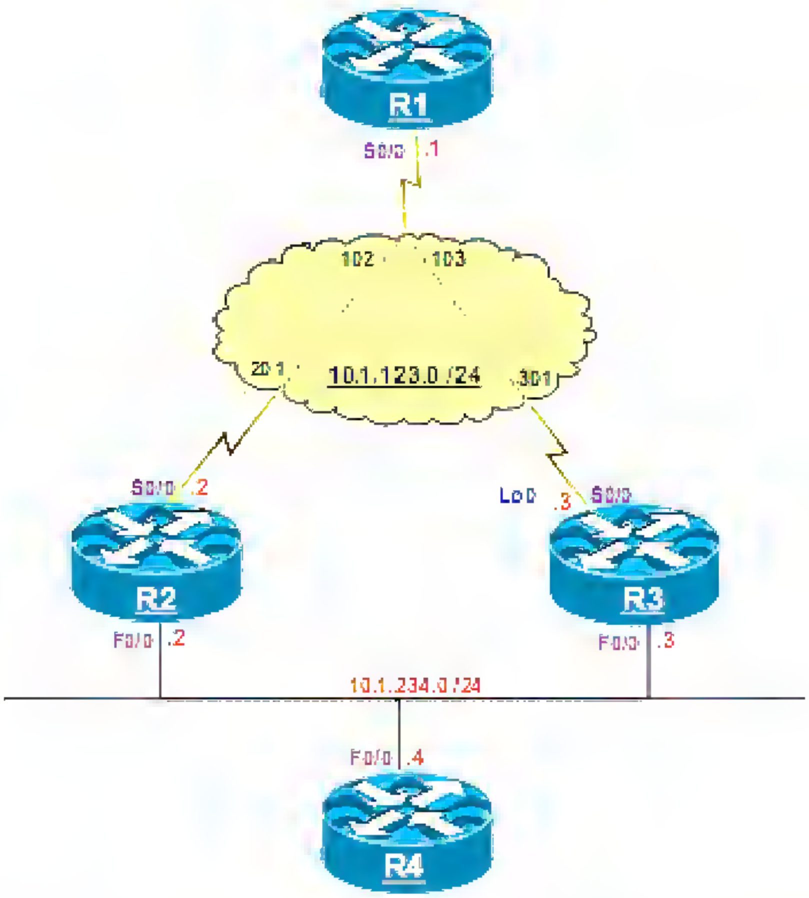

Lab Setup:

- Configure the routers connected to the frame-relay cloud in a hub and spoke manner.
- Router R1 should be configured to be the hub and routers R2 and R3 should be configured as the spokes.
- R1 should be configured directly under the physical interface and it should be configured with two frame-relay mappings, one to R2 and the second one to R3.
- R2 and R3 should be configured directly under their physical interface and they should each be configured with a single frame-relay mapping to the hub.

- The F0/0 interface of R2, R3 and R4 should be configured in VLAN 234.
- Use the following IP addressing chart for IP address assignment.

IP addressing Chart:

Router	Interface / IP addressing
R1	S0/0 = 10.1.123.1 /24
R2	S0/0 = 10.1.123.2 /24 F0/0 = 10.1.234.2 /24
R3	S0/0 = 10.1.123.3 /24 F0/0 = 10.1.234.3 /24
R4	F0/0 = 10.1.234.4 /24

Task 1

Configure HSRP on R2 and R3 using the following requirements:

- R2 should be configured as the active router, whereas, router R3 should be configured as the Standby router
- Preemption should be enabled
- The HSRP IP address should be 10.1.234.100 /24

On R2

R2(config)#**int f0/0**
R2(config-if)#**standby 1 ip 10.1.234.100**
R2(config-if)#**standby 1 priority 110**
R2(config-if)#**standby 1 preempt**

On R3

R3(config)#**int f0/0**
R3(config-if)#**standby 1 ip 10.1.234.100**
R3(config-if)#**standby 1 preempt**

To verify the configuration:

On R3

R3#Show standby

```
FastEthernet0/0 - Group 1
  State is Standby
    1 state change, last state change 00:00:20
  Virtual IP address is 10.1.234.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.107 secs
  Preemption enabled
  Active router is 10.1.234.2, priority 110 (expires in 7.596 sec)
  Standby router is local
  Priority 100 (default 100)
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

On R2

R2#Show standby

```
FastEthernet0/0 - Group 1
  State is Active
    2 state changes, last state change 00:03:59
  Virtual IP address is 10.1.234.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.664 secs
  Preemption enabled
  Active router is local
  Standby router is 10.1.234.3, priority 100 (expires in 7.171 sec)
  Priority 110 (configured 110)
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

Task 2

Configure a default gateway on R4 pointing to the HSRP's IP address. Provide NLRI using RIPv2 on all routers.

On R1, R2 and R3

(config)#router rip
(config-router)#no au
(config-router)#ver 2
(config-router)#network 10.0.0.0

On R4

R4(config)#**ip route** 0.0.0.0 0.0.0.0 10.1.234.100

To verify the configuration:

On R4

R4#Show ip route | b Gateway

```
Gateway of last resort is 10.1.234.100 to network 0.0.0.0

     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.123.0 [120/1] via 10.1.234.3, 00:00:11, FastEthernet0/0
                   [120/1] via 10.1.234.2, 00:00:23, FastEthernet0/0
C       10.1.234.0 is directly connected, FastEthernet0/0
S*   0.0.0.0/0 [1/0] via 10.1.234.100
```

On R1

R1#Show ip route | b Gateway

```
Gateway of last resort is not set

     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.123.0 is directly connected, Serial0/0
R       10.1.234.0 [120/1] via 10.1.123.3, 00:00:09, Serial0/0
                   [120/1] via 10.1.123.2, 00:00:24, Serial0/0
```

Task 3

Configure R2 and R3 to keep track of their Frame-relay connection to R1. This configuration must check end-to-end connectivity, and if R2 fails to reach R1's Frame-relay interface, R3 should become the active router.

IP SLA monitor can be used in conjunction with HSRP. An SLA monitor is configured for testing end-to-end reachability; a track object is configured and tied to SLA monitor. The "track" keyword is used in HSRP to call the tracked object. When connectivity is broken between R1 and R2, the status of the tracked object goes down and HSRP decrements the priority by the configured value, in this case 50 and since the preemption is configured, R3 will generate a Coup message and because it will have a higher priority (100) it will become the active router.

On R2

R2(config)#**ip sla monitor 10**
R2(config-sla-monitor)#**type echo protocol IpIcmpEcho 10.1.123.1 source-interface S0/0**
R2(config-sla-monitor-echo)#**timeout 500**
R2(config-sla-monitor-echo)#**frequency 5**

R2(config)#**Ip SLA monitor schedule 10 start-time now life forever**

R2(config)#**track 1 rtr 1 reachability**

R2(config)#Int F0/0
R2(config-if)#**standby 1 track 1 decrement 50**

On R3

R3(config)#**ip sla monitor 30**
R3(config-sla-monitor)#**type echo protocol IpIcmpEcho 10.1.123.1 source-interface S0/0**
R3(config-sla-monitor-echo)#**timeout 500**
R3(config-sla-monitor-echo)#**frequency 5**

R3(config)#**Ip sla monitor schedule 30 start-time now life forever**

R3(config)#**track 1 rtr 30 reachability**

R3(config)#int f0/0
R3(config-if)#**standby 1 track 1 decrement 50**

To verify the configuration:

On R2

R2#Show standby

```
FastEthernet0/0 - Group 1
 State is Active
  7 state changes, last state change 00:03:57
 Virtual IP address is 10.1.234.100
 Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (v1 default)
 Hello time 3 sec, hold time 10 sec
  Next hello sent in 2.279 secs
 Preemption enabled
 Active router is local
 Standby router is 10.1.234.3, priority 100 (expires in 7.167 sec)
 Priority 110 (configured 110)
  Track object 1 state Up decrement 50
 IP redundancy name is "hsrp-Fa0/0-1" (default)
```

R2#Show ip sla monitor statistics

```
Round trip time (RTT)   Index 1
        Latest RTT: 40 ms
Latest operation start time: *07:28:30.475 UTC Sat Apr 13 2002
Latest operation return code: OK
Number of successes: 50
Number of failures: 0
Operation time to live: Forever
```

On R3

R3#Show standby

```
FastEthernet0/0 - Group 1
 State is Standby
  7 state changes, last state change 00:07:21
 Virtual IP address is 10.1.234.100
 Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (v1 default)
 Hello time 3 sec, hold time 10 sec
  Next hello sent in 0.035 secs
 Preemption enabled
 Active router is 10.1.234.2, priority 110 (expires in 9.142 sec)
 Standby router is local
```

Priority 100 (default 100)
  Track object 1 state Up decrement 50
 IP redundancy name is "hsrp-Fa0/0-1" (default)

R3#Show ip sla monitor statistics

Round trip time (RTT) Index 1
    Latest RTT: 40 ms
Latest operation start time: *12:08:14.903 UTC Fri Apr 12 2002
Latest operation return code: OK
Number of successes: 81
Number of failures: 0
Operation time to live: Forever

To test the configuration:

On R2

R2(config)#int s0/0
R2(config-if)#Shut

You should get the following messages stating that R2 is now in Standby state:

%HSRP-6-STATECHANGE: FastEthernet0/0 Grp 1 state Active -> Speak
%HSRP-6-STATECHANGE: FastEthernet0/0 Grp 1 state Speak -> Standby

To verify the configuration:

On R2

R2#Show standby

FastEthernet0/0 - Group 1
 State is Standby
  9 state changes, last state change 00:01:59
 Virtual IP address is 10.1.234.100
 Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (v1 default)
 Hello time 3 sec, hold time 10 sec
  Next hello sent in 0.347 secs
 Preemption enabled
 Active router is 10.1.234.3, priority 100 (expires in 9.351 sec)
 Standby router is local

```
  Priority 60 (configured 110)
    Track object 1 state Down decrement 50
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

R2#Show track

```
Track 1
  Response Time Reporter 1 reachability
  Reachability is Down
    5 changes, last change 00:03:03
  Latest operation return code: Timeout
  Tracked by:
    HSRP FastEthernet0/0 1
```

On R3

R3#Show standby

```
FastEthernet0/0 - Group 1
  State is Active
    8 state changes, last state change 00:02:31
  Virtual IP address is 10.1.234.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.253 secs
  Preemption enabled
  Active router is local
  Standby router is 10.1.234.2, priority 60 (expires in 9.247 sec)
  Priority 100 (default 100)
    Track object 1 state Up decrement 50
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

To test the configuration further:

On R4

R4#Ping 10.1.123.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.123.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 56/57/61 ms
```

On R2

R2(config-if)#int s0/0
R2(config-if)#no shut

You should receive the following message stating that R2 is once again the active router:

%HSRP-6-STATECHANGE: **FastEthernet0/0 Grp 1 state Standby -> Active**

To test the configuration:

On R4

R4#Ping 10.1.123.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.123.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 56/58/60 ms

Note R4 can successfully Ping R1's Frame-relay interface using the default route that is pointing to the HSRP's IP address of 10.1.234.100

R2#Show standby

```
FastEthernet0/0 - Group 1
  State is Active
    10 state changes, last state change 00:02:23
  Virtual IP address is 10.1.234.100
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.156 secs
  Preemption enabled
  Active router is local
  Standby router is 10.1.234.3, priority 100 (expires in 8.153 sec)
  Priority 110 (configured 110)
    Track object 1 state Up decrement 50
  IP redundancy name is "hsrp-Fa0/0-1" (default)
```

R2#Show track

Track 1

```
 Response Time Reporter 1 reachability
 Reachability is Up
   6 changes, last change 00:04:59
 Latest operation return code: OK
 Latest RTT (millisecs) 40
 Tracked by:
   HSRP FastEthernet0/0 1
```

On R3

R3#Show standby

```
FastEthernet0/0 - Group 1
 State is Standby
   10 state changes, last state change 00:06:07
 Virtual IP address is 10.1.234.100
 Active virtual MAC address is 0000.0c07.ac01
   Local virtual MAC address is 0000.0c07.ac01 (v1 default)
 Hello time 3 sec, hold time 10 sec

    Next hello sent in 1.718 secs
 Preemption enabled
 Active router is 10.1.234.2, priority 110 (expires in 7.709 sec)
 Standby router is local
 Priority 100 (default 100)
   Track object 1 state Up decrement 50
 IP redundancy name is "hsrp-Fa0/0-1" (default)
```

R3#Show track

```
Track 1
 Response Time Reporter 1 reachability
 Reachability is Up
   6 changes, last change 00:23:42
 Latest operation return code: OK
 Latest RTT (millisecs) 39
 Tracked by:
   HSRP FastEthernet0/0 1
```

Task 4

Erase the startup configuration and reload the routers before proceeding to the next lab.

Lab 5 – Object Tracking

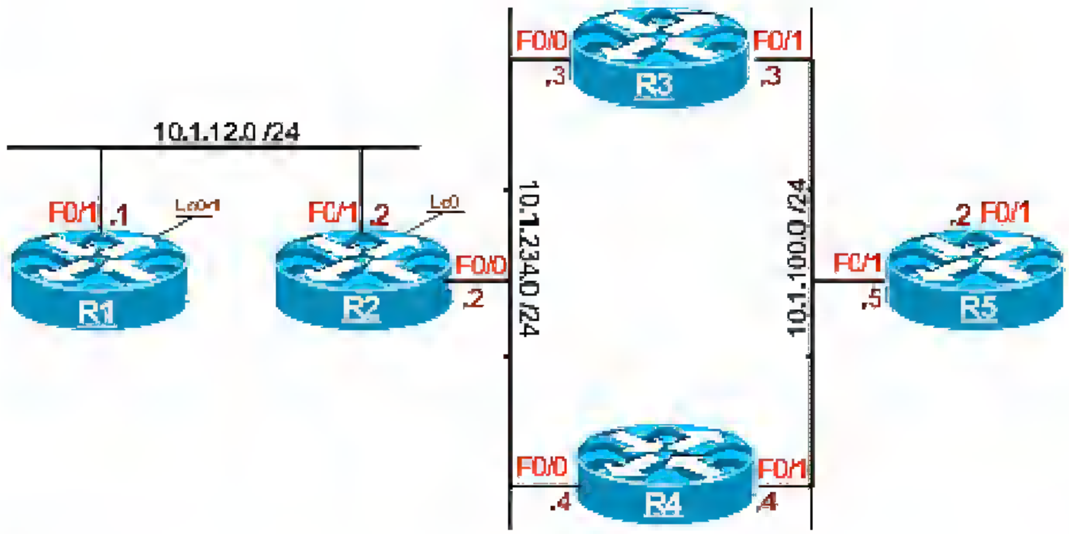

Lab Setup:

- Configure the F0/1 interface of R1 and R2 in Vlan 100
- Configure the F0/0 interface of R2, R3 and R4 in Vlan 200
- Configure the F0/1 interface of R3, R4 and R5 in Vlan 300
- Use the following IP addressing chart:

IP addressing chart:

Router	Interface / IP addressing
R1	F0/1 – 10.1.12.1 /24 Lo0 – 1.1.1.1 /24 Lo1 – 11.1.1.1 /24
R2	F0/1 – 10.1.12.2 /24 F0/0 – 10.1.234.2 Lo0 – 2.2.2.2 /24
R3	F0/0 – 10.1.234.3 /24 F0/1 – 10.1.100.3 /24
R4	F0/0 – 10.1.234.4 /24 F0/1 – 10.1.100.4 /24
R5	F0/1 – 10.1.100.5 /24

Task 1

Configure OSPF area 0 on R1, R2, and on F0/0 and F0/1 interface of R3 and R4. Ensure that all loopback interfaces are advertised with their correct mask.

On R1

```
R1(config-if)#int lo0
R1(config-if)#ip ospf net point-to-point

R1(config-if)#int lo1
R1(config-if)#ip ospf net point-to-point

R1(config)#router ospf 1
R1(config-router)#netw 10.1.12.1 0.0.0.0 are 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 11.1.1.1 0.0.0.0 are 0
```

On R2

```
R2(config-if)#int lo0
R2(config-if)#ip ospf net point-to-point

R2(config-if)#router ospf 1
R2(config-router)#netw 10.1.12.2 0.0.0.0 are 0
R2(config-router)#netw 10.1.234.2 0.0.0.0 are 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 0
```

On R3

```
R3(config-if)#router ospf 1
R3(config-router)#netw 10.1.234.3 0.0.0.0 are 0
R3(config-router)#netw 10.1.100.3 0.0.0.0 are 0
```

On R4

```
R4(config-if)#router ospf 1
R4(config-router)#netw 10.1.234.4 0.0.0.0 are 0
R4(config-router)#netw 10.1.100.4 0.0.0.0 are 0
```

To verify the configuration:

On R4

R4#Show ip route ospf | Inc O

```
O       1.1.1.0 [110/21] via 10.1.234.2, 00:00:46, FastEthernet0/0
O       2.2.2.0 [110/11] via 10.1.234.2, 00:00:46, FastEthernet0/0
O       10.1.12.0 [110/20] via 10.1.234.2, 00:00:46, FastEthernet0/0
O       11.1.1.0 [110/21] via 10.1.234.2, 00:00:46, FastEthernet0/0
```

Task 2

Configure HSRP on F0/1 interface of R3 and R4 using the following policy:

The Virtual IP address – 10.1.100.100
R3 should be the active router
R4 should be the standby router
R4 should become the active router ONLY if R3 is down
R5 should use the active router to reach the networks behind R3 and/or R4

On R3

```
R3(config)#int F0/1
R3(config-if)#standby 1 ip 10.1.100.100
R3(config-if)#standby 1 priority 110
R3(config-if)#standby 1 preempt
```

On R4

```
R4(config)#int F0/1
R4(config-if)#standby 1 ip 10.1.100.100
R4(config-if)#standby 1 preempt
```

To verify the configuration:

On R3

R3#Show standby brief

```
                    P indicates configured to preempt.
                    |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/1        1   110  P Active   local           10.1.100.4      10.1.100.100
```

On R4

R4#Show standby brief

```
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/1        1   100  P Standby 10.1.100.3       local          10.1.100.100
```

To test the configuration:

On R3

```
R3(config)#int F0/1
R3(config-if)#Shut
```

Note R3 is down and the output of the following show command reveals that R4 became the active router:

R4#Show standby brief

```
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/1        1   100  P Active  local           unknown         10.1.100.100
```

On R3

```
R3(config)#int F0/1
R3(config-if)#no shut
```

Note when R3 comes up, it becomes the active router and R4 once again becomes the standby router:

On R3

R3#Show standby brief

```
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active         Standby        Virtual IP
Fa0/1        1   110  P Active  local          10.1.100.4     10.1.100.100
```

On R4

R4#Show standby brief
P indicates configured to preempt.
|

Interface	Grp	Prio	P	State	Active	Standby	Virtual IP
Fa0/1	1	100	P	Standby	10.1.100.3	local	10.1.100.100

The following configures a static default route to provide NLRI to R5:

On R5

```
R5(config)#ip route 0.0.0.0 0.0.0.0 10.1.100.100
```

Task 3

Configure the HSRP routers such that when the F0/0 interface of R3 is down, HSRP is notified after 60 seconds, and R4 becomes the active router; and when R3's F0/0 interface comes back up, HSRP is notified after 2 minutes and R3 resumes the active role. The interface should be polled every 2 seconds for this task.

This task requires the configuration of Enhanced Object Tracking. Before this feature was introduced, HSRP had a simple tracking mechanism that permitted the tracking of a given interface. The Enhanced Object Tracking feature can be used to create a separates tracking process that can be used by other Cisco IOS processes as well as HSRP.
HSRP, VRRP, GLBP and/or IP SLA can be considered as client processes, and these processes can register their interest in tracking objects and be notified when and if the tracked object changes state.
In Enhanced Object Tracking, the tracking process periodically polls the tracked objects in order to detect changes; by default this is done every second and this timer can be changed by using the "Track timer interface *seconds*" global configuration command.
The changes are communicated to the registered client, either immediately or after a configured delay; using the "delay up *seconds*" or "delay down *seconds*" command in track sub-configuration mode.
In this task the tracking process is configured to track the line-protocol state of an interface.

On R3

The following command specifies the interval in which the tracking process polls the tracked objects; this task requires a poll interval of 2 seconds:

R3(config)#track timer interface 2

The following command tracks the line-protocol state of F0/0 interface of R3 and enters the tracking configuration mode:

R3(config)#**track 1 interface F0/0 line-protocol**

The following commands specify a period of time in seconds to delay communicating state changes of a tracked object, the first command instructs the object tracking process to wait for 60 seconds before communicating the DOWN state of the tracked object and the second command instructs the object tracking process to waits for 2 minutes before communicating the UP state of the same tracked object:

R3(config-track)#delay down 60
R3(config-track)#delay up 120

To verify the configuration:

On R3

R3#Show track 1

```
Track 1
  Interface Ethernet0/0 line-protocol
  Line protocol is Up
    3 changes, last change 00:05:04
  Delay up 120 secs, down 60 secs
  Tracked by:
    HSRP FastEthernet0/1 1
```

To test the configuration:

On R3

R3#deb track

Note the interface went down 00:17:41

R3(config)#int F0/0
R3(config-if)#Shut

Note OSPF transitioned from FULL state to down state, because interface was shutdown:

00:17:41.963: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Ethernet0/0 from FULL to DOWN, Neighbor Down: Interface down or detached

00:17:41.971: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.234.4 on FastEthernet0/0 from FULL to DOWN, Neighbor Down: Interface down or detached

The following states that the interface changed state to administratively down:
00:17:43.951: %LINK-5-CHANGED: Interface FastEthernet0/0, changed state to

administratively down Note it took one minute for the process to communicate the

state change
00:17:44.951: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to down

00:18:41: **Track: 1 Down change delay expired**
00:18:41: **Track: 1 Change #4 interface Fa0/0, line-protocol Up->Down**
00:18:41.955: **%TRACKING-5-STATE: 1 interface Fa0/0 line-protocol Up->Down**

Note the following reveals that HSRP transitioned the F0/1 interface of R3 from Active → to Speak → standby:

00:18:42.371: %HSRP-5-STATECHANGE: **FastEthernet0/1 Grp 1 state Active -> Speak**
00:18:52.371: %HSRP-5-STATECHANGE: **FastEthernet0/1 Grp 1 state Speak -> Standby**

To test the second condition:

In this test the F0/0 interface of R3 is brought back up:

R3(config)#int F0/0
R3(config-if)#no shut **Note the interface came up at 00:20:47**

The interface comes up:
00:20:47.327: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
00:20:48.327: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up

OSPF reestablishes the adjacency and transitions into FULL state:

00:20:50.495: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
00:20:50.507: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.234.4 on FastEthernet0/0 from LOADING to FULL, Loading Done

Note the output of the following command reveals that even though R3's interface came back up in UP/UP state, it is still in standby, because the Enhanced Object Tracking has NOT communicated the UP/UP state to its client process in this case HSRP.

```
R3#Show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/1       1    99   P Standby  10.1.100.4      local           10.1.100.100
```

The following show command reveals the count down of the timer.

```
R3#Sh track 1

Track 1
  Interface FastEthernet0/0 line-protocol
  Line protocol is Down, delayed Up (74 secs remaining)
    4 changes, last change 00:02:49
  Delay up 120 secs, down 60 secs
  Tracked by:
    HSRP Ethernet0/1 1

R3#Sh track 1

Track 1
  Interface FastEthernet0/0 line-protocol
  Line protocol is Down, delayed Up (38 secs remaining)
    4 changes, last change 00:03:24
  Delay up 120 secs, down 60 secs
  Tracked by:
    HSRP Ethernet0/1 1

R3#Sh track 1

Track 1
  Interface FastEthernet0/0 line-protocol
  Line protocol is Down, delayed Up (7 secs remaining)
    4 changes, last change 00:03:56
  Delay up 120 secs, down 60 secs
  Tracked by:
    HSRP Ethernet0/1 1
```

The Track timer expires:

```
00:22:45: Track: 1 Up change delay expired
```

Interface F0/0 transitions from Down to UP state:

00:22:45: Track: 1 Change #5 interface Fa0/0, line-protocol Down->Up
00:22:45.343: %TRACKING-5-STATE: 1 interface Fa0/0 line-protocol Down->Up

HSRP detects the change and R3 transitions from Standby to Active state:
00:22:45.435: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Standby -> Active

```
R3#Sh standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/1       1    110  P Active  local           10.1.100.4      10.1.100.100
```

Task 4

Re-configure the HSRP routers such that if Network 1.1.1.0 /24 goes down, R4 becomes the active router. But if Network 1.1.1.0 /24 is up, then, R3 should be the active router.

The first step is to remove the track 1 configuration:

On R3

```
R3(config)#NO track 1
R3(config)#NO track timer interface 2
```

Second step is to configure a new tracking such that Network 1.1.1.0 /24 is tracked in this case reachability is what is tracked:

```
R3(config)#track 1 ip route 1.1.1.0/24 reachability
```

Note the following reveals that R3 is the active router:

```
R3#sh stand brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/1       1    110  P Active  local           10.1.100.4      10.1.100.100
```

To verify the configuration:

On R3

R3#Show track

```
Track 1
  IP route 1.1.1.0 255.255.255.0 reachability
  Reachability is Up (OSPF)
    1 change, last change 00:03:16
  First-hop interface is FastEthernet0/0
  Tracked by:

    HSRP FastEthernet0/1 1
```

To test the configuration:

Note the following summarizes the test procedure:

- **Debug Track is enabled on R3**
- **The network that is being tracked (1.1.1.0 /24) is Shutdown**
- **The debug messages are observed & the change is verified in HSRP**
- **Network 1.1.1.0/24 is brought back up and once again the debug messages are observed and the change in HSRP is verified**

On R3

R3#Debug Track

On R1

```
R1(config)#int lo0
R1(config-if)#Shut
```

On R3

02:37:26: Track: 1 Change #2 IP route 1.1.1.0/24, OSPF->no route, reachability Up->Down
02:37:26.251: %TRACKING-5-STATE: 1 ip route 1.1.1.0/24 reachability Up->Down

02:37:27.243: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Active -> Speak

02:37:37.243: %HSRP-5-STATECHANGE: **FastEthernet0/1 Grp 1 state Speak -> Standby**

R3#Show standby brief
P indicates configured to preempt.
|

Interface	Grp	Prio	P	State	Active	Standby	Virtual IP
Fa0/1	1	99	P	Standby	10.1.100.4	local	10.1.100.100

On R1

R1(config)#int lo0
R1(config-if)#No shut

On R3

02:47:11: **Track: 1 Change #3 IP route 1.1.1.0/24, no route->OSPF, reachability Down->Up**

02:47:11.259: %TRACKING-5-STATE: **1 ip route 1.1.1.0/24 reachability Down->Up**

02:47:12.203: %HSRP-5-STATECHANGE: **FastEthernet0/1 Grp 1 state Standby -> Active**

R3#Show standby brief

P indicates configured to preempt.
|

Interface	Grp	Prio	P	State	Active	Standby	Virtual IP
Fa0/1	1	110	P	Active	local	10.1.100.4	10.1.100.100

Task 5

Configure R3 such that if host 11.1.1.1 /24 is NOT reachable, R4 becomes the active router, but if host 11.1.1.1 /24 is reachable, R3 should remain as the active and R4 should be the standby router. You should use IP SLA with ICMP Echo messages to accomplish this task, the frequency of the messages should be 5 seconds with a timeout of 250 ms.

IP SLA is configured to resolve this task, IP SLA can be configured check network availability: two aspects of an IP SLA operation can be tracked: State & Reachability.

Every IP SLA operation maintains an operation return code, for example: OK, Over

the Threshold and etc. this return code is the return code is interpreted by the tracking process:

Tracking	Return Code	Track State
State	OK All other return code	UP Down
Reachability	OK or Over the Threshold All other return code	UP Down

R3(config)#**ip sla monitor 2**
R3(config-sla-monitor)#**type echo protocol IpIcmpEcho 11.1.1.1**
R3(config-sla-monitor-echo)#**timeout 250**
R3(config-sla-monitor-echo)#**frequency 5**

R3(config)#**ip sla monitor schedule 2 start-time now life forever**

The following command tracks the state of an IP SLA object and enters the tracking configuration mode:

R3(config)#**track 1 rtr 2**

To verify the configuration:

On R3

R3#Sh track 1

```
Track 1
  Response Time Reporter 2 state
  State is Up
    2 changes, last change 00:00:24
  Latest operation return code: OK
  Latest RTT (millisecs) 80
  Tracked by:
    HSRP FastEthernet0/1 1
```

To test the configuration:

Note the following summarizes the test procedure:

- Debug Track is enabled on R3
- The host that is being tracked (11.1.1.0 /24) is Shutdown
- The debug messages are observed & the change is verified in HSRP
- Host 11.1.1.0/24 is brought back up and once again the debug messages are observed and the change in HSRP is verified

On R3

R3#Debug track

On R1

R1(config)#int lo1
R1(config-if)#Shut

On R3

00:52:21: Track: 1 Change #3 rtr 2, state Up->Down
00:52:21.719: %TRACKING-5-STATE: 1 rtr 2 state Up->Down

00:52:24.071: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Active -> Speak

R3#Show standby brie

P indicates configured to preempt.
|

Interface	Grp	Prio	P	State	Active	Standby	Virtual IP
Fa0/1	1	99	P	Standby	10.1.100.4	local	10.1.100.100

On R1

R1(config)#int lo1
R1(config-if)#no shut

On R3

00:56:31: Track: 1 Change #4 rtr 2, state Down->Up

00:56:31.723: %TRACKING-5-STATE: 1 rtr 2 state Down->Up
00:56:33.167: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Standby -> Active

R3#Show standby brie

```
              P indicates configured to preempt.
              |
Interface   Grp  Prio P State    Active          Standby         Virtual IP
Fa0/1       1    110  P Active   local           10.1.100.4      10.1.100.100
```

Task 6

Erase the startup configuration of the routers and reload them before proceeding to the next task.

Advanced
CCIE Routing & Switching
2.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

GRE
Tunnel

Lab 1 – Basic configuration of GRE Tunnel

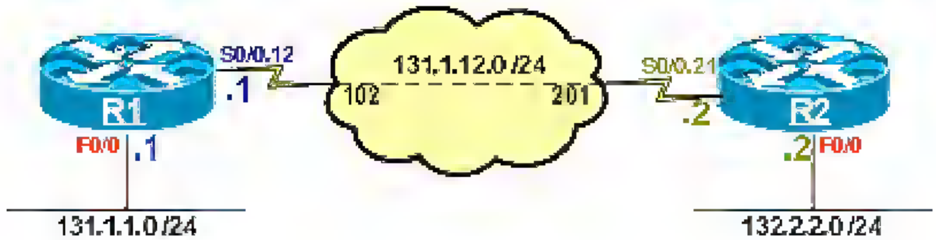

Lab Setup:

- Configure R1 and R2 in a frame-relay point-to-point manner.
- Configure R1's F0/0 interface in VLAN 10.
- Configure R2's F0/0 interface in VLAN 20
- Configure the routers with the IP addressing identified in the above diagram
- R1 should use DLCI 102 to connect to R2 and R2 should use DLCI 201 to connect to R1.

Task 1

Configure OSPF between networks 131.1.1.0 /24 and 132.2.2.0 /24 to provide reachability. You must use a GRE Tunnel and an IP addressing space of your choice to accomplish this task.

On R1

```
R1(config)#int tunnel 12

R1(config-if)#ip address 200.1.12.1 255.255.255.0
R1(config-if)#tunnel source S0/0.12
R1(config-if)#tunnel Destination 131.1.12.2
```

On R2

R2(config)#int tunnel 21
R2(config-if)#ip address 200.1.12.2 255.255.255.0
R2(config-if)#tunnel Source S0/0.21
R2(config-if)#tunnel Destination 131.1.12.1

To test the configuration:

On R2

R2#Ping 200.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms

To configuring OSPF between the networks:

On R2

R2(config)#router ospf 1
R2(config-router)#netw
R2(config-router)#network 200.1.12.2 0.0.0.0 area 0
R2(config-router)#network 132.2.2.2 0.0.0.0 area 0

On R1

R1(config-if)#router ospf 1
R1(config-router)#network 200.1.12.1 0.0.0.0 area 0
R1(config-router)#network 131.1.1.1 0.0.0.0 area 0

To test the configuration:

On R1

R1#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2

```
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    200.1.12.0/24 is directly connected, Tunnel12
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.1.0 is directly connected, FastEthernet0/0
C       131.1.12.0 is directly connected, Serial0/0.12
     132.2.0.0/24 is subnetted, 1 subnets
O       132.2.2.0 [110/11112] via 200.1.12.2, 00:01:01, Tunnel12
```

R1#Ping 132.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 132.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms
```

R1#Show ip ospf neighbor

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
200.1.12.2        0   FULL/  -        00:00:39    200.1.12.2      Tunnel12
```

R1#show ip ospf int tunnel 12

```
Tunnel12 is up, line protocol is up
  Internet Address 200.1.12.1/24, Area 0
  Process ID 1, Router ID 200.1.12.1, Network Type POINT_TO_POINT, Cost: 11111
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

Note, there are no routing protocols running between R1's S0/0/0.12 interface and R2's S0/0/0.21 interface, if OSPF is configured to advertise 131.1.12.0 /24 network, then 131.1.1.0 and 132.2.2.0 networks will be reachable via the serial sub-interface instead of the tunnel interface. If OSPF is running on network 131.1.12.0 /24, one way to force the traffic to go through the tunnel interface instead of the serial sub-interface is to assign a higher OSPF cost to the serial sub-interface or assign a lower OSPF cost to the tunnel interface, but this has to be configured on both routers, as follows:

On R2

R2(config)#router ospf 1
R2(config-router)#network 131.1.12.2 0.0.0.0 area 0

R2#Show ip ospf int tunnel 12

Tunnel21 is up, line protocol is up
Internet Address 200.1.12.1/24, Area 0
Process ID 1, Router ID 200.1.12.1, Network Type POINT_TO_POINT, Cost: 11111
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

R2#Show ip ospf int S0/0.21

Serial0/0.21 is up, line protocol is up
Internet Address 131.1.12.2/24, Area 0
Process ID 1, Router ID 200.1.12.2, Network Type POINT_TO_POINT, Cost: 64
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

R2(config)#int tunnel 21
R2(config-if)#**ip ospf cost 63**

On R1

R1(config)#router ospf 1
R1(config-router)#network 131.1.12.1 0.0.0.0 area 0

R1#Show ip ospf int tunnel 12

Tunnel12 is up, line protocol is up
Internet Address 200.1.12.1/24, Area 0
Process ID 1, Router ID 200.1.12.1, Network Type POINT_TO_POINT, Cost: 11111
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

R1#Show ip ospf int S0/0.12

Serial0/0.12 is up, line protocol is up
Internet Address 131.1.12.1/24, Area 0
Process ID 1, Router ID 200.1.12.1, Network Type POINT_TO_POINT, Cost: 64

Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

```
R1(config)#int tunnel 12
R1(config-if)#ip ospf cost 63

R1#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    200.1.12.0/24 is directly connected, Tunnel12
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.1.0 is directly connected, FastEthernet0/0
C       131.1.12.0 is directly connected, Serial0/0.12
     132.2.0.0/24 is subnetted, 1 subnets
O       132.2.2.0 [110/64] via 200.1.12.2, 00:00:13, Tunnel12
```

Task 2

Re-configure the routers and use the IP addressing specified as per the following diagram. Ensure that failure of any of the links will not bring down the tunnel. Use the following policy to accomplish this task:

- Use the same IP address space for the tunnel that you used in the previous task
- Run a routing protocol of your choice to accomplish this task
- R1 and R2's F0/1 interface must be configured in VLAN 12.
- R1's F0/0 interface should be configured in VLAN 11 and R2's F0/0 interface should be configured in VLAN 22
- Create Loopback 0 interfaces on the routers, R1's Lo0 should be 1.1.1.1 /8 and R2's Lo0 should be configured to be 2.2.2.2 /8
- Run OSPF between the 131.1.1.0 /24 and 132.2.2.0 /24 networks and ensure that this traffic between these networks uses the tunnel interface.

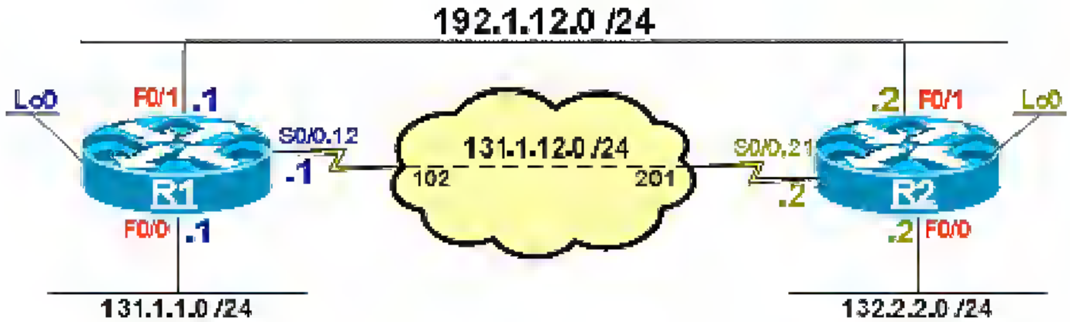

In this task you are asked to run a routing protocol of your choice, in order to provide reachability, this configuration will demo and configure RIPv2, Eigrp and then OSPF, the first option is to run RIPv2:

On R1

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#network 131.1.0.0
R1(config-router)#network 192.1.12.0
R1(config-router)#passive-interface F0/0
R1(config-router)#network 1.0.0.0

On R2

R2(config)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#network 192.1.12.0
R2(config-router)#network 131.1.0.0
R2(config-router)#network 2.0.0.0

On R1

R1#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

```
     ia - IS-IS inter area, * - candidate default, U - per-user static route
     o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.1.12.0/24 is directly connected, FastEthernet0/1
C    1.0.0.0/8 is directly connected, Loopback0
R    2.0.0.0/8 [120/1] via 192.1.12.2, 00:00:05, FastEthernet0/1
               [120/1] via 131.1.12.2, 00:00:02, Serial0/0.12
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.1.0 is directly connected, FastEthernet0/0
C       131.1.12.0 is directly connected, Serial0/0.12
```

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
     D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
     N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
     E1 - OSPF external type 1, E2 - OSPF external type 2
     i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
     ia - IS-IS inter area, * - candidate default, U - per-user static route
     o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.1.12.0/24 is directly connected, FastEthernet0/1
R    1.0.0.0/8 [120/1] via 192.1.12.1, 00:00:11, FastEthernet0/1
               [120/1] via 131.1.12.1, 00:00:08, Serial0/0.21
C    2.0.0.0/8 is directly connected, Loopback0
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.12.0 is directly connected, Serial0/0.21
     132.2.0.0/24 is subnetted, 1 subnets
C       132.2.2.0 is directly connected, FastEthernet0/0
```

Next step is to create the tunnel interface:

On R1

```
R1(config)#int tunnel 12
R1(config-if)#ip address 200.1.12.1 255.255.255.0
R1(config-if)#tunnel source lo0
R1(config-if)#tunnel destination 2.2.2.2
```

```
R1(config)#router ospf 1
R1(config-router)#network 200.1.12.1 0.0.0.0 area 0
R1(config-router)#network 131.1.1.1 0.0.0.0 area 0
```

On R2

```
R2(config)#int tunnel 21
R2(config-if)#ip address 200.1.12.2 255.255.255.0
R2(config-if)#tunnel source lo0
R2(config-if)#tunnel destination 1.1.1.1

R2(config)#router ospf 1
R2(config-router)#network 200.1.12.2 0.0.0.0 area 0
R2(config-router)#network 132.2.2.2 0.0.0.0 area 0
```

On R1

```
R1#Show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.1.12.0/24 is directly connected, FastEthernet0/1
C    1.0.0.0/8 is directly connected, Loopback0
R    2.0.0.0/8 [120/1] via 192.1.12.2, 00:00:11, FastEthernet0/1
               [120/1] via 131.1.12.2, 00:00:11, Serial0/0.12
C    200.1.12.0/24 is directly connected, Tunnel12
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.1.0 is directly connected, FastEthernet0/0
C       131.1.12.0 is directly connected, Serial0/0.12
     132.2.0.0/24 is subnetted, 1 subnets
O       132.2.2.0 [110/11112] via 200.1.12.2, 00:07:18, Tunnel12
```

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.1.12.0/24 is directly connected, FastEthernet0/1
R    1.0.0.0/8 [120/1] via 192.1.12.1, 00:00:14, FastEthernet0/1
               [120/1] via 131.1.12.1, 00:00:15, Serial0/0.21
C    2.0.0.0/8 is directly connected, Loopback0
C    200.1.12.0/24 is directly connected, Tunnel21
     131.1.0.0/24 is subnetted, 2 subnets
O       131.1.1.0 [110/11112] via 200.1.12.1, 00:08:08, Tunnel21
C       131.1.12.0 is directly connected, Serial0/0.21
     132.2.0.0/24 is subnetted, 1 subnets
C       132.2.2.0 is directly connected, FastEthernet0/0
```

Testing Eigrp:

On R1

```
R1(config)#router eigrp 100
R1(config-router)#no au
R1(config-router)#network 1.0.0.0
R1(config-router)#network 131.1.12.0 0.0.0.255
R1(config-router)#network 192.1.12.0
R1(config-router)#exit
R1(config)#no router rip
```

On R2

```
R2(config)#router eigrp 100
R2(config-router)#no au

R2(config-router)#network 2.0.0.0
R2(config-router)#network 192.1.12.0
R2(config-router)#network 131.1.12.0 0.0.0.255
```

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.1.12.0/24 is directly connected, FastEthernet0/1
C    1.0.0.0/8 is directly connected, Loopback0
D    2.0.0.0/8 [90/156160] via 192.1.12.2, 00:01:00, FastEthernet0/1
C    200.1.12.0/24 is directly connected, Tunnel12
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.1.0 is directly connected, FastEthernet0/0
C       131.1.12.0 is directly connected, Serial0/0.12
     132.2.0.0/24 is subnetted, 1 subnets
O       132.2.2.0 [110/11112] via 200.1.12.2, 00:00:49, Tunnel12
```

To test OSPF:

On R1

```
R1(config)#router ospf 1
R1(config-router)#network 131.1.12.1 0.0.0.0 area 0
R1(config-router)#network 192.1.12.1 0.0.0.0 area 0
R1(config-router)#network 1.1.1.1 0.0.0.0 area 0

R1(config)# NO router eigrp 100
```

On R2

```
R2(config)#router ospf 1
R2(config-router)#network 2.2.2.2 0.0.0.0 area 0
R2(config-router)#network 131.1.12.2 0.0.0.0 area 0
R2(config-router)#network 192.1.12.2 0.0.0.0 area 0

R2(config-router)#NO router eigrp 100
```

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    192.1.12.0/24 is directly connected, FastEthernet0/1
C    1.0.0.0/8 is directly connected, Loopback0
     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 192.1.12.2, 00:18:33, FastEthernet0/1
C    200.1.12.0/24 is directly connected, Tunnel12
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.1.0 is directly connected, FastEthernet0/0
C       131.1.12.0 is directly connected, Serial0/0.12
     132.2.0.0/24 is subnetted, 1 subnets
O       132.2.2.0 [110/2] via 192.1.12.2, 00:18:34, FastEthernet0/1
```

Note 132.2.2.0 /24 network is now reachable via F0/1, but the policy of this task states that this network should be reachable via the Tunnel interface. One way to resolve this problem is to use PBR, as follows:

On R1

R1(config)#**access-list** 100 **permit ip** 131.1.1.0 0.0.0.255 132.2.2.0 0.0.0.255

R1(config)#route-map **TEST** permit 10
R1(config-route-map)#match ip addr 100
R1(config-route-map)#set interface tunnel 12

R1(config)#route-map **TEST** permit 20

R1(config)#ip local policy route-map TEST

R1(config)#int f0/0
R1(config-if)#ip policy route-map TEST

To test the configuration:

On R1

R1#Traceroute

```
Protocol [ip]:
Target IP address: 132.2.2.2
Source address: 131.1.1.1
Numeric display [n]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 132.2.2.2

  1 200.1.12.2 0 msec *  0 msec
```

Note the traffic traverses through the tunnel interface

The same configuration must be performed on R2.

Task 3

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 2 – Configuration of GRE Tunnel II

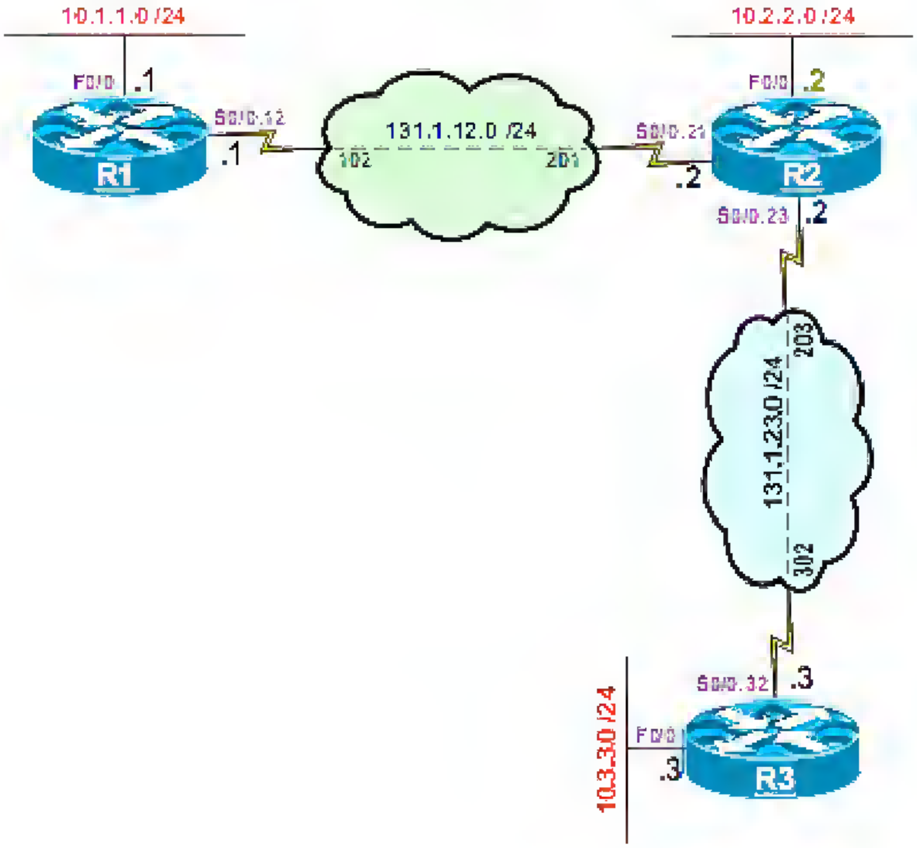

Lab Setup:

- Configure the frame-relay connection between R1, R2 and R2 and R3 in a point-to-point manner.
- Configure R1's F0/0 interface in VLAN 10.
- Configure R2's F0/0 interface in VLAN 20.
- Configure R3's F0/0 interface in VLAN 30
- Configure the routers with the IP addressing identified in the above diagram

Task 1

Configure OSPF between the Private networks (10.1.1.0, 10.2.2.0 and 10.3.3.0 /24) and ensure that these networks have reachability to each other. You must configure GRE Tunnel interface/s to accomplish this task. DO NOT run a routing protocol, Static routes or PBR for the links connecting R1 to R2 and R2 to R3.

On R1

```
R1(config)#int tunnel 12
R1(config-if)#ip addr 200.1.12.1 255.255.255.0
R1(config-if)#tunnel source S0/0.12
R1(config-if)#tunnel destination 131.1.12.2
```

On R2

```
R2(config)#int tunnel 21
R2(config-if)#ip addr 200.1.12.2 255.255.255.0
R2(config-if)#tunnel source S0/0.21
R2(config-if)#tunnel destination 131.1.12.1

R2(config)#int tunnel 23
R2(config-if)#ip address 200.1.23.2 255.255.255.0
R2(config-if)#tunnel source S0/0.23
R2(config-if)#tunnel destination 131.1.23.3
```

On R3

```
R3(config)#int tunnel 32
R3(config-if)#ip address 200.1.23.3 255.255.255.0
R3(config-if)#tunnel source S0/0.32
R3(config-if)#tunnel destination 131.1.23.2
```

To test the configuration:

On R3

```
R3#Ping 200.1.23.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.23.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms
```

On R2

R2#Ping 200.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms

R2#Ping 200.1.23.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.23.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/76 ms

On R1

R1#Ping 200.1.12.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms

To configure OSPF routing protocol on the private networks:

On R1

R1(config)#router ospf 1
R1(config-router)#network 200.1.12.1 0.0.0.0 area 0
R1(config-router)#network 10.1.1.1 0.0.0.0 area 0

On R2

R2(config)#router ospf 1
R2(config-router)#network 200.1.12.2 0.0.0.0 area 0
R2(config-router)#network 200.1.23.2 0.0.0.0 area 0
R2(config-router)#network 10.2.2.2 0.0.0.0 area 0

On R3

R3(config)#router ospf 1
R3(config-router)#network 200.1.23.3 0.0.0.0 area 0

R3(config-router)#network 10.3.3.3 0.0.0.0 area 0

To test the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

O    200.1.23.0/24 [110/22222] via 200.1.12.2, 00:00:28, Tunnel12
C    200.1.12.0/24 is directly connected, Tunnel12
     10.0.0.0/24 is subnetted, 3 subnets
O       10.3.3.0 [110/22223] via 200.1.12.2, 00:00:28, Tunnel12
O       10.2.2.0 [110/11112] via 200.1.12.2, 00:00:28, Tunnel12
C       10.1.1.0 is directly connected, FastEthernet0/0
     131.1.0.0/24 is subnetted, 1 subnets
C       131.1.12.0 is directly connected, Serial0/0.12
```

R1#Ping 10.2.2.2

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms
```

R1#Ping 10.3.3.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 140/140/140 ms
```

Task 2

Remove the configuration from the previous task and re-configure the routers based on the Lab setup guide in the beginning of this lab.

On R1

R1(config)#NO router ospf 1
R1(config)#NO interface Tunnel 12

On R2

R2(config)#NO router ospf 1
R2(config)#NO interface Tunnel 21
R2(config)#NO interface Tunnel 23

On R3

R3(config)#NO router ospf 1
R3(config)#NO interface Tunnel 32

Task 3

Create the following Loopback 0 interfaces and establish the tunnel based on these interfaces, run a routing protocol of your choice and DO NOT configure static routes to accomplish this task.
R1's Lo0 = 1.1.1.1 /8, R2's Lo0 = 2.2.2.2 /8 and R3's Lo0 = 3.3.3.3 /8

On R1

R1(config)#int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0

R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#network 131.1.0.0
R1(config-router)#network 1.0.0.0

On R2

R2(config)#int lo0
R2(config-if)#ip address 2.2.2.2 255.0.0.0

R2(config-if)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#network 131.1.0.0
R2(config-router)#network 2.0.0.0

On R3

R3(config)#int lo0
R3(config-if)#ip address 3.3.3.3 255.0.0.0

R3(config-if)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#network 131.1.0.0
R3(config-router)#netw 3.0.0.0

To test the configuration:

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    1.0.0.0/8 [120/2] via 131.1.23.2, 00:00:01, Serial0/0.32
R    2.0.0.0/8 [120/1] via 131.1.23.2, 00:00:01, Serial0/0.32
C    3.0.0.0/8 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.3.3.0 is directly connected, FastEthernet0/0
     131.1.0.0/24 is subnetted, 2 subnets
R       131.1.12.0 [120/1] via 131.1.23.2, 00:00:01, Serial0/0.32
C       131.1.23.0 is directly connected, Serial0/0.32
```

To configure the Tunnel:

On R1

```
R1(config)#int tunnel 12
R1(config-if)#ip address 200.1.12.1 255.255.255.0
R1(config-if)#tunnel source Lo0
R1(config-if)#tunnel destination 2.2.2.2
```

On R2

```
R2(config)#int tunnel 21
R2(config-if)#ip address 200.1.12.2 255.255.255.0
R2(config-if)#tunnel source Lo0
R2(config-if)#tunnel destination 1.1.1.1

R2(config)#int tunnel 23
R2(config-if)#ip address 200.1.23.2 255.255.255.0
R2(config-if)#tunnel source lo0
R2(config-if)#tunnel destination 3.3.3.3
```

On R3

```
R3(config)#int tunnel 32
R3(config-if)#ip address 200.1.23.3 255.255.255.0
R3(config-if)#tunnel source lo0
R3(config-if)#tunnel destination 2.2.2.2
```

To test the configuration:

On R1

```
R1#Ping 200.1.12.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms
```

On R2

```
R2#Ping 200.1.12.1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 200.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms

R2#Ping 200.1.23.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.23.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms

On R3

R3#Ping 200.1.23.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.23.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 72/72/72 ms

To configure OSPF on the routers:

On R1

R1(config)#router ospf 1
R1(config-router)#network 10.1.1.1 0.0.0.0 area 0
R1(config-router)#network 200.1.12.1 0.0.0.0 area 0

On R2

R2(config)#router ospf 1
R2(config-router)#network 10.2.2.2 0.0.0.0 area 0
R2(config-router)#network 200.1.12.2 0.0.0.0 area 0
R2(config-router)#network 200.1.23.2 0.0.0.0 area 0

On R3

R3(config)#router ospf 1
R3(config-router)#network 10.3.3.3 0.0.0.0 area 0
R3(config-router)#network 200.1.23.3 0.0.0.0 area 0

To test the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
R    2.0.0.0/8 [120/1] via 131.1.12.2, 00:00:17, Serial0/0.12
O    200.1.23.0/24 [110/22222] via 200.1.12.2, 00:04:38, Tunnel12
R    3.0.0.0/8 [120/2] via 131.1.12.2, 00:00:17, Serial0/0.12
C    200.1.12.0/24 is directly connected, Tunnel12
     10.0.0.0/24 is subnetted, 3 subnets
O       10.3.3.0 [110/22223] via 200.1.12.2, 00:04:38, Tunnel12
O       10.2.2.0 [110/11112] via 200.1.12.2, 00:04:38, Tunnel12
C       10.1.1.0 is directly connected, FastEthernet0/0
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.12.0 is directly connected, Serial0/0.12
R       131.1.23.0 [120/1] via 131.1.12.2, 00:00:18, Serial0/0.12
```

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    1.0.0.0/8 [120/1] via 131.1.12.1, 00:00:21, Serial0/0.21
```

```
C    2.0.0.0/8 is directly connected, Loopback0
C    200.1.23.0/24 is directly connected, Tunnel23
R    3.0.0.0/8 [120/1] via 131.1.23.3, 00:00:18, Serial0/0.23
C    200.1.12.0/24 is directly connected, Tunnel21
     10.0.0.0/24 is subnetted, 3 subnets
O       10.3.3.0 [110/11112] via 200.1.23.3, 00:03:43, Tunnel23
C       10.2.2.0 is directly connected, FastEthernet0/0
O       10.1.1.0 [110/11112] via 200.1.12.1, 00:03:43, Tunnel21
     131.1.0.0/24 is subnetted, 2 subnets
C       131.1.12.0 is directly connected, Serial0/0.21
C       131.1.23.0 is directly connected, Serial0/0.23
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    1.0.0.0/8 [120/2] via 131.1.23.2, 00:00:19, Serial0/0.32
R    2.0.0.0/8 [120/1] via 131.1.23.2, 00:00:19, Serial0/0.32
C    200.1.23.0/24 is directly connected, Tunnel32
C    3.0.0.0/8 is directly connected, Loopback0
O    200.1.12.0/24 [110/22222] via 200.1.23.2, 00:02:40, Tunnel32
     10.0.0.0/24 is subnetted, 3 subnets
C       10.3.3.0 is directly connected, FastEthernet0/0
O       10.2.2.0 [110/11112] via 200.1.23.2, 00:02:40, Tunnel32
O       10.1.1.0 [110/22223] via 200.1.23.2, 00:02:40, Tunnel32
     131.1.0.0/24 is subnetted, 2 subnets
R       131.1.12.0 [120/1] via 131.1.23.2, 00:00:20, Serial0/0.32
C       131.1.23.0 is directly connected, Serial0/0.32
```

Task 4

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 3 – Configuration of GRE Tunnel III

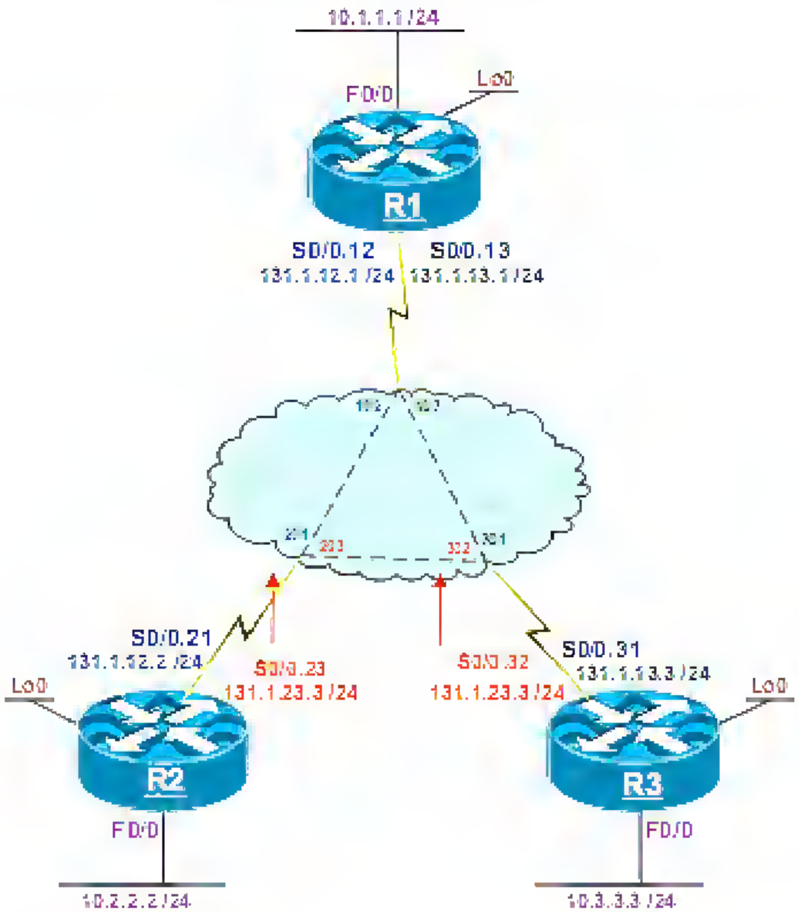

Lab Setup:

- The routers should be configured in a full mesh point-to-point Frame-relay
- The F0/0 interface of R1 should be in VLAN 10, R2's F0/0 should be configured in VLAN 20 and R3's F0/0 should be configured in VLAN 30.
- Use the following IP addressing chart to configure IP addressing on the routers.
- Run RIPv2 on the links that connect the routers (131.1.12.0, 131.1.13.0, 131.1.23.0).

IP addressing and DLCI information Chart:

Routers / Intf	IP address	Local DLCI
R1 S0/0.12 R1 S0/0.13 F0/0 Lo0	131.1.12.1 /24 131.1.13.1 /24 10.1.1.1 /24 1.1.1.1 /8	102 103
R2 S0/0.21 R2 S0/0.23 F0/0 Lo0	131.1.12.2 /24 131.1.23.2 /24 10.2.2.2 /24 2.2.2.2 /8	201 203
R3 S0/0.31 R3 S0/0.32 F0/0 Lo0	131.1.13.3 /24 131.1.23.2 /24 10.3.3.3 /24 3.3.3.3 /8	301 302

Task 1

Run OSPF on the Private networks and ensure that these networks can reach each other through the GRE Tunnel. When configuring the Tunnel interface/s ensure that redundancy is provided such that if a link is down the tunnel remains in up state and reachability is provided through the alternate route.

To accomplish this task the most reliable interface should be used to establish the tunnels, remember that you should have reachability to the Loopback interfaces via multiple paths or else the redundancy will not work.

On R1

```
R1(config)#router rip
R1(config-router)#network 1.0.0.0
R1(config-router)#network 131.1.0.0
```

On R2

```
R2(config)#router rip
R2(config-router)#network 2.0.0.0
R2(config-router)#network 131.1.0.0
```

On R3

```
R3(config)#router rip
```

R3(config-router)#network 3.0.0.0
R3(config-router)#network 131.1.0.0

To establish the tunnels:

On R1

R1(config)#int tunnel 12
R1(config-if)#ip address 200.1.12.1 255.255.255.0
R1(config-if)#tunnel source lo0
R1(config-if)#tunnel destination 2.2.2.2

R1(config)#int tunnel 13
R1(config-if)#ip address 200.1.13.1 255.255.255.0
R1(config-if)#tunnel source Lo0
R1(config-if)#tunnel destination 3.3.3.3

On R2

R2(config)#int tunnel 21
R2(config-if)#ip address 200.1.12.2 255.255
R2(config-if)#tunnel source Lo0
R2(config-if)#tunnel destination 1.1.1.1

R2(config)#int tunnel 23
R2(config-if)#ip address 200.1.23.2 255.255
R2(config-if)#tunnel source Lo0
R2(config-if)#tunnel destination 3.3.3.3

On R3

R3(config)#int tunnel 31
R3(config-if)#ip address 200.1.13.3 255.255.255.0
R3(config-if)#tunnel source Lo0
R3(config-if)#tunnel destination 1.1.1.1

R3(config)#int tunnel 32
R3(config-if)#ip address 200.1.23.3 255.255.255.0
R3(config-if)#tunnel source Lo0
R3(config-if)#tunnel destination 2.2.2.2

To run OSPF in the tunnel:

On R1

R1(config)#router ospf 1
R1(config-router)#network 200.1.12.1 0.0.0.0 area 0
R1(config-router)#network 200.1.13.1 0.0.0.0 area 0
R1(config-router)#network 10.1.1.1 0.0.0.0 area 0

On R2

R2(config)#router ospf 1
R2(config-router)#network 10.2.2.2 0.0.0.0 area 0
R2(config-router)#network 200.1.12.2 0.0.0.0 area 0
R2(config-router)#network 200.1.23.2 0.0.0.0 area 0

On R3

R3(config)#router ospf 1
R3(config-router)#network 10.3.3.3 0.0.0.0 area 0
R3(config-router)#network 200.1.13.3 0.0.0.0 area 0
R3(config-router)#network 200.1.23.3 0.0.0.0 area 0

To test the configuration:

On R1

R1#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C    1.0.0.0/8 is directly connected, Loopback0
R    2.0.0.0/8 [120/1] via 131.1.12.2, 00:00:25, Serial0/0.12
O    200.1.23.0/24 [110/22222] via 200.1.13.3, 00:00:20, Tunnel13
                   [110/22222] via 200.1.12.2, 00:00:20, Tunnel12
R    3.0.0.0/8 [120/1] via 131.1.13.3, 00:00:03, Serial0/0.13
C    200.1.12.0/24 is directly connected, Tunnel12
```

```
C    200.1.13.0/24 is directly connected, Tunnel13
     10.0.0.0/24 is subnetted, 3 subnets
O       10.3.3.0 [110/11112] via 200.1.13.3, 00:00:20, Tunnel13
O       10.2.2.0 [110/11112] via 200.1.12.2, 00:00:21, Tunnel12
C       10.1.1.0 is directly connected, FastEthernet0/0
     131.1.0.0/24 is subnetted, 3 subnets
C       131.1.12.0 is directly connected, Serial0/0.12
C       131.1.13.0 is directly connected, Serial0/0.13
R       131.1.23.0 [120/1] via 131.1.13.3, 00:00:05, Serial0/0.13
                   [120/1] via 131.1.12.2, 00:00:27, Serial0/0.12
```

On R2

R2#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    1.0.0.0/8 [120/1] via 131.1.12.1, 00:00:04, Serial0/0.21
C    2.0.0.0/8 is directly connected, Loopback0
C    200.1.23.0/24 is directly connected, Tunnel23
R    3.0.0.0/8 [120/1] via 131.1.23.3, 00:00:26, Serial0/0.23
C    200.1.12.0/24 is directly connected, Tunnel21
O    200.1.13.0/24 [110/22222] via 200.1.23.3, 00:03:16, Tunnel23
                   [110/22222] via 200.1.12.1, 00:03:16, Tunnel21
     10.0.0.0/24 is subnetted, 3 subnets
O       10.3.3.0 [110/11112] via 200.1.23.3, 00:03:16, Tunnel23
C       10.2.2.0 is directly connected, FastEthernet0/0
O       10.1.1.0 [110/11112] via 200.1.12.1, 00:03:17, Tunnel21
     131.1.0.0/24 is subnetted, 3 subnets
C       131.1.12.0 is directly connected, Serial0/0.21
R       131.1.13.0 [120/1] via 131.1.23.3, 00:00:01, Serial0/0.23
                   [120/1] via 131.1.12.1, 00:00:06, Serial0/0.21
C       131.1.23.0 is directly connected, Serial0/0.23
```

On R3

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    1.0.0.0/8 [120/1] via 131.1.13.1, 00:00:09, Serial0/0.31
R    2.0.0.0/8 [120/1] via 131.1.23.2, 00:00:08, Serial0/0.32
C    200.1.23.0/24 is directly connected, Tunnel32
C    3.0.0.0/8 is directly connected, Loopback0
O    200.1.12.0/24 [110/22222] via 200.1.23.2, 00:04:46, Tunnel32
                   [110/22222] via 200.1.13.1, 00:04:46, Tunnel31
C    200.1.13.0/24 is directly connected, Tunnel31
     10.0.0.0/24 is subnetted, 3 subnets
C       10.3.3.0 is directly connected, FastEthernet0/0
O       10.2.2.0 [110/11112] via 200.1.23.2, 00:04:47, Tunnel32
O       10.1.1.0 [110/11112] via 200.1.13.1, 00:04:47, Tunnel31
     131.1.0.0/24 is subnetted, 3 subnets
R       131.1.12.0 [120/1] via 131.1.23.2, 00:00:09, Serial0/0.32
                   [120/1] via 131.1.13.1, 00:00:10, Serial0/0.31
C       131.1.13.0 is directly connected, Serial0/0.31
C       131.1.23.0 is directly connected, Serial0/0.32
```

To test the configuration:

On R3

```
R3(config)#int S0/0.32
R3(config-subif)#shut
```

R3#Show ip route

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
      o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

R    1.0.0.0/8 [120/1] via 131.1.13.1, 00:00:11, Serial0/0.31
R    2.0.0.0/8 [120/2] via 131.1.13.1, 00:00:11, Serial0/0.31
C    200.1.23.0/24 is directly connected, Tunnel32
C    3.0.0.0/8 is directly connected, Loopback0
O    200.1.12.0/24 [110/22222] via 200.1.13.1, 00:00:00, Tunnel31
C    200.1.13.0/24 is directly connected, Tunnel31
     10.0.0.0/24 is subnetted, 3 subnets
C       10.3.3.0 is directly connected, FastEthernet0/0
O       10.2.2.0 [110/22223] via 200.1.13.1, 00:00:00, Tunnel31
O       10.1.1.0 [110/11112] via 200.1.13.1, 00:00:01, Tunnel31
     131.1.0.0/24 is subnetted, 3 subnets
R       131.1.12.0 [120/1] via 131.1.13.1, 00:00:12, Serial0/0.31
C       131.1.13.0 is directly connected, Serial0/0.31
R       131.1.23.0 [120/2] via 131.1.13.1, 00:00:14, Serial0/0.31
```

Note R3's routing table converged; therefore, R3 can reach 10.1.1.0 /24 and 10.2.2.0 /24 through 200.1.13.1.

Task 2

Erase the startup config and reload the routers before proceeding to the next lab.

Lab 4 – GRE & Recursive loops

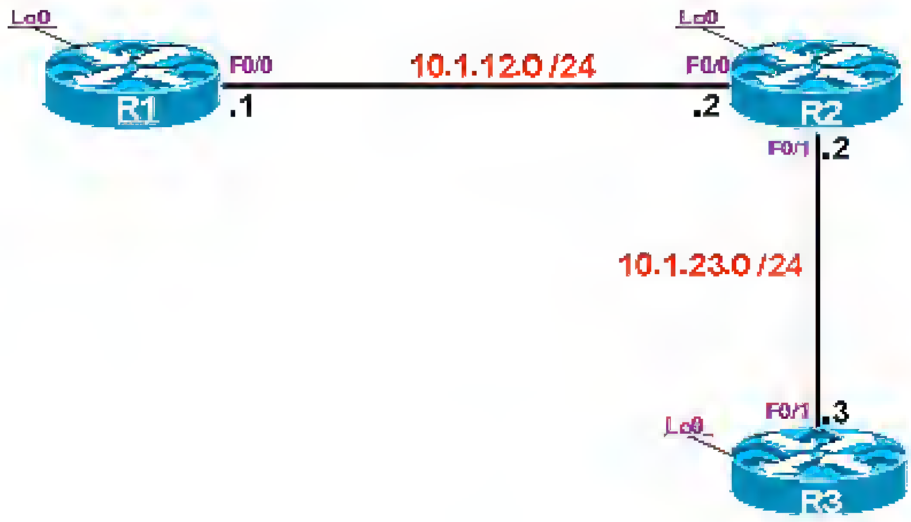

Lab Setup:

- Configure the F0/0 interface of R1 and R2 in VLAN 12.
- Configure the F0/1 interface of R2, R3 in VLAN 23.

IP addressing:

Router	Interface / IP address
R1	F0/0 = 10.1.12.1 /24 Loopback0 = 1.1.1.1 /8
R2	F0/0 = 10.1.12.2 /24 F0/1 = 10.1.23.2 /24 Loopback0 = 2.2.2.2 /8

R3	F0/1 = 10.1.23.3 /24 Loopback0 = 3.3.3.3 /8

Task 1

Configure RIPv2 on these routers and advertise their directly connected networks.

On R1

```
R1(config)#router rip
R1(config-router)#no au
R1(config-router)#ver 2
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 10.0.0.0
```

On R2

```
R2(config)#router rip
R2(config-router)#no au
R2(config-router)#ver 2
R2(config-router)#netw 10.0.0.0
```

On R3

```
R3(config)#router rip
R3(config-router)#no au
R3(config-router)#ver 2
R3(config-router)#netw 10.0.0.0
R3(config-router)#netw 3.0.0.0
```

To verify the configuration:

On R1

R1#Sh ip route rip

```
R    3.0.0.0/8 [120/2] via 10.1.12.2, 00:00:14, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.23.0 [120/1] via 10.1.12.2, 00:00:14, FastEthernet0/0
```

On R2

R2#Show ip route rip

```
R    1.0.0.0/8 [120/1] via 10.1.12.1, 00:00:24, Ethernet0/0
R    3.0.0.0/8 [120/1] via 10.1.23.3, 00:00:02, Ethernet0/1
```

On R3

R3#Show ip route rip

```
R    1.0.0.0/8 [120/2] via 10.1.23.2, 00:00:13, FastEthernet0/1
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.12.0 [120/1] via 10.1.23.2, 00:00:13, FastEthernet0/1
```

Task 2

Configure a GRE tunnel from R1 to R3, the ip address of this tunnel interface should be 200.1.1.1 /24 and 200.1.1.3 for R1 and R3 respectively. The tunnel source should be based on the loopback0 interface of these routers. R3 should use Lo0 interface of R1, whereas, R1 should use Lo0 interface of R3 as their tunnel destination.

On R1

```
R1(config)#int tu1
R1(config-if)#ip addr 200.1.1.1 255.255.255.0
R1(config-if)#tunnel source lo0
R1(config-if)#tunnel destination 3.3.3.3
```

On R3

```
R3(config)#int tu1
R3(config-if)#ip addr 200.1.1.3 255.255.255.0
R3(config-if)#tunnel source lo0
R3(config-if)#tunnel destination 1.1.1.1
```

To verify and test the configuration:

On R1

R1#Show ip int brie | Inc Tunnel1

```
Tunnel1                200.1.1.1      YES manual up                    up
```

R1#Ping 200.1.1.3

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/73/100 ms
```

On R3

R3#Show ip int brie | Inc Tunnel1

```
Tunnel1                200.1.1.3      YES manual up                    up
```

R3#Ping 200.1.1.1

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/74/136 ms
```

Note the Tunnel1 interface is in UP/UP state.

Task 3

Configure Eigrp 100 through the tunnel interface of R1 and R3; these routers should advertise their Loopback0 interface in this routing protocol.

On R1

```
R1(config)#Router eigrp 100
R1(config-router)#no au
R1(config-router)#netw 200.1.1.0
R1(config-router)#netw 1.0.0.0
```

On R3

R3(config)#Router eigrp 100
R3(config-router)#no au
R3(config-router)#netw 200.1.1.0
R3(config-router)#netw 3.0.0.0

Note you should get the following messages on these routers, the following shows the messages received on R1:

%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 200.1.1.3 (Tunnel1) is up: new adjacency
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 200.1.1.3 (Tunnel1) is down: holding time expired
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 200.1.1.3 (Tunnel1) is up: new adjacency

%TUN-5-RECURDOWN: Tunnel1 temporarily disabled due to recursive routing
*Mar 1 01:14:54.243: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state to down
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 200.1.1.3 (Tunnel1) is down: interface down

The tunnel interface status depends on the IP reachability to the tunnel destination; in this case R1 and R3 find that reachability through RIPv2. But since the tunnel destination (Lo0) is also advertised through Eigrp, therefore, when the tunnel comes up, these routers find a better route through Eigrp to the tunnel destination (Eigrp's Administrative distance being lower than RIP), therefore, they find the reachability to the tunnel destination through the tunnel and that's what causes the recursive loops.

Task 4

Configure R1 and R3 such that the tunnel interface does NOT flap and it stays in UP/UP state, you should NOT configure a static route, or stop advertising this interface in Eigrp to accomplish this task.

To fix this problem an access-list list is configured to identify the IP address of the tunnel destination, then, the administrative distance of this route is configured to be lower than Eigrp's administrative distance.
Note once these commands are entered, the routers will find a better route (RIPv2) than Eigrp to the IP address of the tunnel destination.

On R1

R1(config)#access-list 1 permit 3.0.0.0

R1(config)#Router rip
R1(config-router)#distance 89 0.0.0.0 255.255.255.255 1

```
R1#Clear ip route *

R1#Sh ip route rip

R    3.0.0.0/8 [89/2] via 10.1.12.2, 00:00:02, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.23.0 [120/1] via 10.1.12.2, 00:00:02, FastEthernet0/0
```

Task 5

Erase the startup configuration and reload the routers before proceeding to the next lab.